

Toward Empathic Systems

Implicit Understanding and Modulation of
Human Cognitive and Affective States

Héctor López Carral

TESI DOCTORAL UPF / year 2022

THESIS SUPERVISOR

Prof. Dr. Paul F.M.J. Verschure

Department Departament de Tecnologies de la Informació i
les Comunicacions



“The question is not whether intelligent machines can have any emotions, but whether machines can be intelligent without emotions”
—Marvin Minsky

Thanks

First of all, I would like to thank my supervisor, Paul Verschure, for giving me the opportunity to do this PhD and all the work contained here.

Thanks to all the SPECS members that have been involved, one way or another, in these years at the lab. Thanks to Pedro Omedas and Riccardo Zucca for inviting me into SPECS in the first place and guiding me even before I started this PhD. Thanks to Anna Mura for all your help. Thanks to Diogo Santos-Pata, Klaudia Grechuta, and Giovanni Maffei for your support and for showing me the way, in different senses. Thanks to Maria Blancas, Sock Ching Low, Martina Maier, Jordi-Ysard Puigbó, Vicky Vouloutsi, Dina Urich, Belén Rubio, and Sytse Wierenga for all the learning and experiences.

Thanks to more recent members of SPECS, with whom I have shared this journey or part of it: Ismael Freire, Adrián Fernández, Óscar Guerrero, Francisco Santos, Marco Galli, Albert Albesa, and so many others. Thanks to Quique Martínez and Alex Escuredo for all the technical support. Thanks to Àdria España and to Xavi de la Torre, for too many things.

Thanks to my parents and my sister for always supporting me and having enabled all of my education leading me up to this point. Thanks to Clothilde for giving me so much support and understanding over this last year. And thanks to so many other people in my life, currently present or not.

Thanks to Lydia Garcia for guiding me throughout this process from the academic and administrative side. Thanks to Boris Bellalta for your support in this last period.

Thanks to so many artists for having provided the music that has fueled so much of this work.

I would also like to thank IBEC and the Spanish government for having economically supported this work through my fellowship.

I could not have done this without so many people, mentioned here or not, so I can only feel immense gratitude toward everyone that is or has been part of my life. I thank you all.

Abstract

Emotions are a fundamental part of human life. They play a critical role in how we think, behave, and implicitly understand each other. However, although computing devices are increasingly consequential in our world, they still lack meaningful emotional capabilities, so we can only interact with them explicitly. This thesis explores and develops how synthetic systems could implicitly understand and modulate human cognitive and affective states to enhance our interaction with them. First, we introduce a technological architecture to sense user information, interpret it, and dynamically adapt immersive environments. Then, we present several studies exploring methods to infer different internal states from multiple sources, including physiological signals, keystroke dynamics, and affective ratings. Finally, we show two examples of interactive and adaptive experiences exploiting implicit understanding to assist users when needed. Overall, our results offer insights into human emotion and contribute toward developing empathic systems, better prepared to support us and to act autonomously.

Resum

Les emocions són una part fonamental de la vida humana. Tenen un paper crític en com pensem, ens comportem i ens entenem implícitament. Tanmateix, tot i que els dispositius informàtics són cada cop més importants al nostre món, encara no tenen capacitats emocionals significatives, de manera que només podem interactuar amb ells de manera explícita. Aquesta tesi explora i desenvolupa com els sistemes sintètics podrien comprendre i modular implícitament els estats cognitius i afectius humans per millorar la nostra interacció amb ells. En primer lloc, introduïm una arquitectura tecnològica per detectar informació de l'usuari, interpretar-la i adaptar dinàmicament entorns immersius. A continuació, presentem diversos estudis que exploren mètodes per inferir diferents estats interns a partir de múltiples fonts, incloent senyals fisiològics, dinàmiques de pulsació de tecles i valoracions afectives. Finalment, mostrem dos exemples d'experiències interactives i adaptatives que utilitzen la comprensió implícita per ajudar els usuaris quan sigui necessari. Els nostres resultats ofereixen nous coneixements sobre l'emoció humana i contribueixen a desenvolupar sistemes empàtics, més preparats per donar-nos suport i actuar de manera autònoma.

Contents

List of figures	xviii
List of tables	xix
Publications	xix
1 INTRODUCTION	1
1.1 Emotions	3
1.1.1 Theories of Emotion	5
1.1.2 Models of Emotion	7
1.1.3 Neurophysiology of Emotion	10
1.1.4 Measuring Emotion	14
1.2 Human-Computer Interaction	20
1.3 Affective Computing	23
1.4 Thesis Outline	25
I Architecture for Implicit Understanding of Internal States	31
2 A DISTRIBUTED SENSING ARCHITECTURE FOR PSYCHOPHYSIOLOGICAL EXPERIMENTATION IN VIRTUAL AND MIXED REALITY	33
2.1 Introduction	34
2.1.1 State of the Art	35

2.2	Architecture	36
2.2.1	Data Transmission and Synchronization	37
2.2.2	Sensing Devices	38
2.2.3	Online Analysis	40
2.2.4	Recording and Offline Analysis	41
2.2.5	Immersive Environment	42
2.3	Application	43
2.3.1	Sensing Devices and Signals	44
2.3.2	Experimental Environment	44
2.3.3	Online Analysis	45
2.3.4	Offline Storage and Analysis	46
2.4	Performance	46
2.5	Potential Applications	48
2.6	Conclusion	48
3	A VIRTUAL REALITY SYSTEM FOR THE SIMULATION OF NEURODIVERSITY	51
3.1	Introduction	52
3.1.1	Autism Spectrum Disorder	52
3.1.2	Virtual Reality for Neurodiversity Simulation	53
3.2	The Neurodiversity Experience	56
3.2.1	Stimuli	56
3.2.2	Implementation	56
3.3	Discussion and Conclusions	60
3.3.1	Further Steps	60
4	A SOCIALLY COOPERATIVE COGNITIVE ARCHITECTURE FOR REHABILITATION	63
4.1	Introduction	64
4.1.1	Conceptual Framework	66
4.2	Architecture	69
4.2.1	Layered Structure	70
4.2.2	Digital Twin	72
4.2.3	Serious Gaming Scenarios	75

4.3 Discussion	76
--------------------------	----

II Methods for Implicit Understanding of Internal States **79**

5 INDUCTION OF UNCONSCIOUS PROCESSING OF EMOTIONAL STIMULI IN AN IMMERSIVE SYSTEM	81
5.1 Introduction	82
5.1.1 Implicit User States in Human-Computer Interaction	82
5.1.2 Unconscious Stimuli Processing	83
5.1.3 Subliminal Stimulation in an Ecologically Valid Navigation Task	84
5.2 Materials and Methods	85
5.2.1 Experiment 1	86
5.2.2 Experiment 2	91
5.3 Results	94
5.3.1 Experiment 1	94
5.3.2 Experiment 2	97
5.4 Discussion and Conclusion	100
6 KEYSTROKE DYNAMICS CORRELATE WITH AFFECTIVE CONTENT	103
6.1 Introduction	104
6.2 Methods	106
6.2.1 Participants	106
6.2.2 Experimental Protocol	106
6.2.3 Data Collection and Processing	107
6.3 Results	109
6.4 Discussion and Conclusion	112
7 SUBJECTIVE RATINGS OF EMOTIVE STIMULI PREDICT THE IMPACT OF THE COVID-19 QUARANTINE ON	

AFFECTIVE STATES	117
7.1 Introduction	118
7.2 Materials and methods	120
7.2.1 Participants	120
7.2.2 Materials	121
7.2.3 Procedure	124
7.2.4 Data analysis	125
7.3 Results	125
7.4 Discussion	130
8 IMPLICIT MEASURES IDENTIFY DIFFERENTIAL EF-	
ECTS OF THE COVID-19 QUARANTINE ON THE PSY-	
CHOLOGICAL HEALTH OF DIVERSE POPULATIONS	137
8.1 Introduction	138
8.2 Methods	140
8.2.1 Participants	140
8.2.2 Materials	140
8.2.3 Procedure	143
8.2.4 Implicit Interaction Measures	143
8.3 Results	144
8.4 Discussion	147
9 BINARY AFFECTIVE RATINGS THROUGH SWIPING	151
9.1 Introduction	152
9.2 Methodology	155
9.3 Results	158
9.4 Discussion	160
III Interactive and Adaptive Systems	163
10 AN ASSISTIVE SYSTEM FOR TRANSFERRING DOMAIN	
KNOWLEDGE TO NOVICE USERS	165
10.1 Introduction	166
10.1.1 Introduction to recommender systems	167

10.1.2	Recommender System for Domain Knowledge Transfer	169
10.2	Use Case: Investigation Knowledge Graph Exploration	172
10.3	Discussion and Conclusion	175
11	EXPERIMENTAL EVALUATION OF AN IMMERSIVE, INTERACTIVE AND CONFLUENT SYSTEM	179
11.1	Introduction	180
11.2	Immersion Evaluation	183
11.2.1	Background	183
11.2.2	Objective	184
11.2.3	Method	184
11.2.4	Results	188
11.2.5	Discussion	190
11.3	Explicit Interaction Evaluation	191
11.3.1	Background	191
11.3.2	Objective	192
11.3.3	Method	193
11.3.4	Results	195
11.3.5	Discussion	197
11.4	Implicit Interaction Evaluation	197
11.4.1	Background	198
11.4.2	Objective	200
11.4.3	Method	200
11.4.4	Results	210
11.4.5	Discussion	213
11.5	Conclusions	214
IV	General Discussion and Conclusions	219
12	GENERAL DISCUSSION	221
13	GENERAL CONCLUSIONS	227

List of Figures

- 1.1 Ekman’s six basic emotions 8
- 1.2 Circumplex model of affect 10
- 1.3 Emotional pathways in the human brain 12
- 1.4 Reality-Virtuality Continuum 22

- 2.1 Implemented Distributed Sensing Architecture (DSA) as used in an experiment 37
- 2.2 Schematic illustration of the eXperience Induction Machine (XIM) 42
- 2.3 Plot of synchronized signals during an experiment 47

- 3.1 Screenshot of the classroom environment 53
- 3.2 Setup of the experience 58
- 3.3 Architecture to process the physiological signals to estimate arousal levels online for interaction with the virtual reality experience 59

- 4.1 Diagram representing the general cognitive architecture based on the Distributed Adaptive Control (DAC) framework 68
- 4.2 Schematic view of the socially cooperative cognitive architecture 69
- 4.3 Digital Twin architecture 73

- 5.1 Schematic illustration of the XIM and EEG montages used. 85

5.2	Stimulus sequence timeline showing the two types of possible trials in Experiment 1	89
5.3	Stimulus sequence timeline showing the two types of possible trials in Experiment 2	93
5.4	Behavioral results from Experiment 1	95
5.5	Behavioral results from Experiment 2	98
6.1	The selection process of the 46 images included in the experimental dataset	108
6.2	Visualization of keystrokes during an example image description	109
6.3	Correlation analysis of each of the extracted features with valence and arousal considering bins of these affective ratings	111
6.4	Correlation analysis of the mean value of the extracted features for each image with valence and arousal	113
7.1	Example of digital assessments of the arousal and pleasure using the Affective Slider	122
7.2	Distribution of the valence and arousal rating for the 30 images selected for this study and the OASIS data set of 900 images	123
7.3	Linear correlations between the ratings obtained in our study and those from OASIS	126
7.4	Shift in the affective ratings for neutral and positive images	127
7.5	Correlations between the differences in valence ratings per participant and self-reported situation during the quarantine period	128
7.6	Differences in valence ratings relative to those of OASIS between participants who, during the quarantine, lived alone and those who did not	129
8.1	Example of usage of the Affective Slider in our study	141
8.2	Images selected from the Open Affective Standardized Image Set (OASIS)	141

8.3	Linear correlations between the ratings obtained in our study and those from OASIS	144
8.4	Shift in the affective ratings for neutral images	145
8.5	Correlation between mouse speed and mean valence rating for each image	146
8.6	Analysis of the text typed by the participants describing the best and the worst things that happened to them during the then-ongoing COVID-19 quarantine period	147
9.1	Affective swiping app (ASA) and the method applied to rate the images	156
9.2	Subset of images selected from OASIS	157
9.3	Relationship between the percentage of like choices and continuous valence	159
9.4	Implicit measures correlate with absolute valence polarity	159
10.1	Abstract conceptualization of the cognitive architecture of the knowledge transfer system based on the DAC framework.	170
10.2	The user interface of the knowledge graph exploration task	174
10.3	Example of a suggested node	176
11.1	Setups used in the Immersion Evaluation study	186
11.2	3D visual representation of the Connectome dataset used in the current study	187
11.3	Differences between the High and Low Immersion conditions in the Immersion Evaluation	190
11.4	Differences between the Explicit Interaction and No Interaction conditions in the Explicit Interaction Evaluation	196
11.5	Conceptual representation of a user interacting in the XIM	203
11.6	Screenshots of starting and ending points in the ‘hard’ level of graph complexity at the start and at the end of the task	204
11.7	Overview of the types of SA interventions triggered by different user states	209

11.8	Visualization of the number of times the SA was triggered by an increase or decrease in arousal and cognitive load by level of difficulty	210
11.9	Differences between the Control, Congruent, and Incongruent conditions in the Implicit Interaction Evaluation .	217

List of Tables

- 1.1 Categories of affective computing 24
- 2.1 Signals captured and transmitted during an experiment using the DSA 45
- 3.1 Examples of stimuli used in the experience, divided between auditory and visual 56
- 6.1 Features extracted from participant’s behavior while typing 110
- 11.1 Experimental design of the Implicit Interaction Evaluation 205

Publications

Included in this thesis

López-Carral, H., Santos-Pata, D., Zucca, R., and Verschure, P. F. (2019). How you type is what you type: Keystroke dynamics correlate with affective content. *2019 8th International Conference on Affective Computing and Intelligent Interaction, ACII 2019*, pages 359–363

López-Carral, H., Grechuta, K., and Verschure, P. F. (2020). Subjective ratings of emotive stimuli predict the impact of the COVID-19 quarantine on affective states. *PLoS ONE*, 15(8 August):1–15

López-Carral, H., Blancas-Muñoz, M., Mura, A., Omedas, P., España-Cumellas, À., Martínez-Bueno, E., Milliken, N., Moore, P., Haque, L., Gilroy, S., and Verschure, P. F. M. J. (2022). A virtual reality system for the simulation of neurodiversity. In Yang, X.-S., Sherratt, S., Dey, N., and Joshi, A., editors, *Proceedings of Sixth International Congress on Information and Communication Technology*, pages 523–531, Singapore. Springer Singapore

López-Carral, H. and Verschure, P. F. (2022). An Assistive System for Transferring Domain Knowledge to Novice Officers. *European Law Enforcement Research Bulletin*, 22(6)

López-Carral, H., Omedas, P., Zucca, R., and Verschure, P. F. (2023d).

A distributed sensing architecture for psychophysiological experimentation in virtual reality. *Manuscript in preparation*

López-Carral, H., de la Torre Costa, J., Freire, I. T., Mura, A., and Verschure, P. F. (2023a). A socially cooperative cognitive architecture for rehabilitation. *Manuscript in preparation*

López-Carral, H., Zucca, R., Omedas, P., and Verschure, P. F. (2023e). Induction of unconscious processing of emotional stimuli in an immersive system. *Manuscript in preparation*

López-Carral, H. and Verschure, P. F. (2023). Binary affective ratings through swiping. *Manuscript in preparation*

López-Carral, H., Omedas, P., and Verschure, P. F. (2023c). Experimental evaluation of an immersive, interactive and confluent system. *International Journal of Human-Computer Studies (submitted)*

López-Carral, H., Grechuta, K., and Verschure, P. F. (2023b). Implicit measures identify differential effects of COVID-19 on the psychological health of diverse populations. *Manuscript in preparation*

Other publications

Santos-Pata, D., Zucca, R., López-Carral, H., and Verschure, P. F. (2019). Modulating grid cell scale and intrinsic frequencies via slow high-threshold conductances: A simplified model. *Neural Networks*

Araya León, M. J., López-Carral, H., Mura, A., and Verschure, P. F. (2023). Environmental and emotional well-being in the spatial experience in the restaurant context, from a work perspective. *Manuscript in preparation*

Chapter 1

INTRODUCTION

We live in a world in constant change due to technological advances. Computing devices are increasingly ubiquitous in our lives and we rely on them for a wide variety of tasks, from working to connecting with other people. Furthermore, our healthcare is also increasingly dependent on the latest technological advancements paired with novel scientific research. Because of this, the need for interacting with these devices as naturally as possible is also increasing. Traditionally, human-computer interaction has relied on the explicit input of the user via devices such as keyboards, mice, or touchscreens. However, more recent advancements in fields such as not only computing, but also neuroscience and psychology, have started to allow the possibility of also taking into account implicit input: the internal states of the users, including their cognitive or affective states (Andre, 2013; Wagner et al., 2013; Serim and Jacucci, 2019; Flavell et al., 2022).

Computers are commonly seen as logical devices that do not have emotional capabilities, and for good reason. However, researchers are working to change this by developing computers that might be able to perceive and express emotions. While most consumer computers currently do not have these capabilities, it is an area of active research. Over the last few decades, the field of affective computing has formally established itself, researching how computers could sense the feelings that their

users are experiencing, and even how these devices could simulate affective states of their own, as well as the consequences that such possibilities would bring over our lives (Picard, 1997).

The motivation behind this endeavor comes from the fact that emotions play a fundamental role in human experience, from cognition to perception, as demonstrated by a multitude of scientific research (Dolan, 2002; Barrett et al., 2016). This suggests that bringing affective capabilities to computers would make them more suitable in assisting humans by optimizing the interaction process, as well as making them more autonomous devices capable of taking decisions that might more closely resemble those that a human would take, by including more implicit and explicit factors (Picard, 1995).

The possibilities of optimizing human-computer interaction by providing affective capabilities to computers are particularly relevant not only for our personal devices in everyday life but also for high-demanding, professional computing tasks. A relevant example of this would be the interaction with large amounts of complex data, such as massive neurological datasets. Optimizing the interaction process with this data would thus help in solving the so-called data deluge phenomenon, in which finding meaning in large and complex datasets is becoming increasingly challenging (Wagner et al., 2013).

Other modern computing possibilities to improve human-computer interaction for certain tasks are the usage of mixed or virtual reality. These interaction paradigms, although envisioned and established several decades ago, have begun to be more broadly available only in recent years, thanks to the advancement in technological capabilities. Mixed and virtual reality systems allow for a deeper coupling between the user and the digital content, which can aid in the visualization and exploration of multidimensional and relational datasets (Lessiter et al., 2011), for example.

In recent years, with the technological advancements in areas such as wearable devices and machine learning, coupled with scientific progress in neuroscience and psychology, the possibility of monitoring an array of signals from individuals during their daily activities has materialized

(Johnson and Picard, 2020; Cosoli et al., 2021). This increases the need for methods to meaningfully extract insights from this data to have an impact on people’s lives. This would be the aim of empathic systems: achieving a deeper understanding of users by leveraging different sources of information provided implicitly that reflect what the user feels. This could have implications not only in enhancing the human-computer interaction process to make it more natural but also in key areas such as digital health, for example, by being able to monitor the mental well-being of people who might be at risk of developing an emotional disorder.

This thesis presents a series of advancements to extend the capabilities of interactive systems to implicitly understand and modulate human cognitive and affective states. In order to develop this project, multidisciplinary integration of knowledge from different areas was necessary. Because of this, a brief survey of the state of the art is provided in the next chapter. We will start by discussing emotions: what they are, how they are generated considering both the body, the brain, and the mind; how they affect us, how we can model them, and how we can measure them. Then, we will pivot towards human-computer interaction, seeing how the field has evolved since its inception, its current trends, and its future direction. Combining what we have seen in the previous two sections, we will discuss the field of affective computing, describing not only its goal, methods, and applications, but also its limitations.

1.1 Emotions

Emotions are a fundamental part of the human experience. Because of this, it is a topic that has been thoroughly discussed throughout human history with gradually evolving interpretations. However, although there has been significant progress in research on emotion in fields such as psychology and neuroscience, there is still no clear consensus on the exact nature and functional role of emotions, as we will see in this section.

Nowadays, when dealing with the wide topic of emotions in science, different terms are used, including “emotion”, “mood”, and “affect”. Al-

though occasionally these terms were used interchangeably, important distinctions have been established. “Affect” is used as a broad term that encompasses both “emotion” and “mood”, as the experiential state of feeling. As such, affective states vary in multiple aspects, such as intensity, specificity, or duration. An “emotion” is a short-duration affective state, generally in response to a specific external stimulus. Meanwhile, a “mood” is an affective state that is maintained for longer than an emotion, is less intense, and is not necessarily caused by a specific external source (Niven, 2013; Ekkekakis, 2013).

Another key term for this thesis is “empathy”. We can define empathy as the capacity for an individual to understand or feel what another individual is experiencing (Stocks and Lishner, 2012; Davis, 2022). Although popularly this might be focused on emotions, it also covers cognitive aspects (Elliott et al., 2011; Riess, 2017). This distinction between emotional or affective empathy and cognitive empathy exists at both the clinical and the neural level (Cox et al., 2012; Shamay-Tsoory et al., 2009). Despite this separation, recent multidimensional approaches have emphasized the necessity of including both (Davis, 2022). As empathy plays a key role in interpersonal interaction, it has also been shown to have significant importance in prosocial behaviors (Stocks and Lishner, 2012; Davis, 2022; Riess, 2017). A deficit of empathy is linked to several psychological disorders (Davis, 2022), including autism spectrum disorder, borderline personality disorder, narcissistic personality disorder, and psychopathy (Rinaldi et al., 2021).

Despite the lack of consensus and diversity of theories on many aspects surrounding emotions, it is now generally accepted that they play a relevant role in many cognitive processes, including memory, perception, and decision-making (Barrett et al., 2016). Furthermore, it is generally agreed that emotions involve subjective perception, physiological responses, and expressive behavior (Gross and Barrett, 2011), although the latter has been more debated due to individual and cultural differences (Barrett, 2018).

In the next subsections, we will give an overview of the theories of emotion, analyzing how the concept of emotion has evolved throughout

history, the most influential theories, and the current outlook. Then, we will see two of the most important models of emotions that are widely used nowadays. Next, we will focus on how emotions are processed and generated in the brain, looking at the known pathways and key areas. Finally, we will offer a broad survey on the diversity of methods that exist to measure emotions and their correlates, in terms of physiology, behavior, and self-assessments.

1.1.1 Theories of Emotion

Classical views of emotions place them as antagonists of rationality, as exemplified by Greek philosophers. In the works of Aristotle, he characterized emotions (using the term *pathos*) as responses to our environment that are linked to pleasure or pain and are similar to desire or appetite, playing an intrinsic part in our lives (Schmitter, 2016). On the other hand, stoic philosophers considered these passions as something to avoid in the pursuit of reason and a virtuous life (Schmitter, 2016).

This general view prevailed until the study of emotions started to become more scientific towards the end of the 19th century, starting with the 1872 book by Charles Darwin *The Expression of the Emotions in Man and Animals* (Darwin, 1872). In his work, Darwin studied emotions from an evolutionary perspective and their origin in animal behavior, studying expressions of emotions, with emphasis on human facial expressions, and their social value as a communicative medium.

A few years later, in 1884, William James published his book *What is an Emotion?* (James, 1884), in which he linked bodily responses to stimuli and the subjective experience of emotion. Around the same time, a physiologist named Carl Lange independently proposed a related idea (Barrett, 2017b). For this reason, we now know this early theory of emotion as the James-Lange theory, a name assigned by philosopher John Dewey. It is worth noting that, according to recent investigations, James did not claim that each emotion has a distinct bodily reaction, only each *instance* of emotion (Barrett, 2017b). Nevertheless, the so-called James-Lange theory is based on the idea that an external stimulus leads to a phys-

iological reaction, which is then interpreted by the brain as an emotion. Subsequently, Walter Cannon and Philip Bard disagreed with this theory, arguing that the conscious feeling of an emotion happens at the same time as the physiological reaction (Cannon, 1927). This is now known as the Cannon-Bard theory.

After the James-Lange and the Cannon-Bard theories were proposed, the study of emotions gained significance, leading to the development of many different theories. In general, these can be broadly categorized depending on where the emphasis is placed: are emotions mostly a physical or cognitive event? Do emotions start in the body or in the brain?

Richard Lazarus argued that, in an emotional instance, the first thing that happens is a cognitive appraisal, which then triggers a series of physiological changes, leading to a behavioral response (Lazarus et al., 1970). Since then, other researchers have proposed influential theories based on this underlying idea (Frijda et al., 1986; Scherer, 1999; Moors et al., 2013).

On the other hand, Antonio Damasio proposed the “somatic marker hypothesis” in his popular book *Descartes’ Error* (Damasio, 1994). According to this hypothesis, each emotion is grounded on a unique physiological fingerprint that is interpreted by the brain to guide our actions in response to the stimulus that caused them (Damasio, 1996). In his book, he also argued that emotion is necessary for rational decision-making and that emotions should be studied considering both the body and the mind (arguing against René Descartes’ dualism). In more recent years, Damasio reviewed his theory, decreasing the importance of body sensing areas of the cortex in favor of subcortical circuits that receive primary sensory signals from the body (Damasio and Carvalho, 2013; LeDoux and Brown, 2017).

As of today, the scientific debate on whether emotions are initiated as a physiological response or as a subjective mental feeling continues. A popular theory proposed by psychologist Lisa Feldman Barrett in recent years claims that emotions are constructed predictively in the brain, integrating internal signals from the body (interoception) and using past experiences (which are also mediated by culture) (Barrett, 2017a,b). This

theory, claims that the properties of affect (arousal and valence, based on the circumplex model of emotion defined by psychologist James Russell (Russell, 1980)) are basic features of consciousness Barrett (2017b). In accordance with other theories and authors, it establishes that the brain regulates the organism through allostasis: a process of regulation to dynamically maintain all internal variables within optimal ranges based on the motivation, goals, and drives of an individual (Barrett, 2017b; Vouloutsi and Verschure, 2018). Therefore, affect plays a role in regulating behavior, as a direct consequence of allostasis (Barrett, 2017b; Vouloutsi and Verschure, 2018).

1.1.2 Models of Emotion

Models of emotion can be broadly grouped into being categorical or dimensional (Sreeja and Mahalakshmi, 2017). Categorical models are based on discrete categories labeled emotions, as defined culturally. Meanwhile, dimensional models rely on the definition of a series of continuous dimensions, with each specific emotion felt being a particular combination of them. Here, we will focus on the two main models used nowadays and discussed throughout this thesis.

The most classical categorical model was established by psychologist Paul Ekman, defining six basic emotions: anger, disgust, fear, happiness, sadness, and surprise (Ekman, 1992) (Figure 1.1). He defined these as separate and discrete emotional states, which differ from each other in meaningful ways to make them distinguishable. Importantly, all basic emotions also have common characteristics, such as rapid onset, short duration, and automatic appraisal. According to this theory, these emotions are universally expressed and recognized, without the need for cultural learning, as they are “preprogrammed” and involuntary (although also affected by individual life experiences). Other emotions would be specific cases or combinations of these basic emotions, although more basic emotions could also exist, as Ekman discussed later, with “contempt” being added as a seventh basic emotion (Ekman and Cordaro, 2011).

Among the dimensional models, one of the most popular ones is the




Figure 1.1: Ekman's six basic emotions

Photos of actors representing the six basic emotions proposed by Paul Ekman. From the top left: anger, fear, disgust, surprise, happiness, and sadness. Original images from Ekman and Friesen (1976).

circumplex model of emotions, initially proposed by psychologist James Russell in 1980 (Russell, 1980). This model uses two dimensions commonly defined as arousal (from deactivation to activation) and valence (from unpleasant to pleasant, also called "pleasure"). Other variations include additional dimensions, such as "tension" in Wundt's original model from 1897 (Wundt, 1897), or "dominance" in Mehrabian's model (Mehrabian, 1980). However, arousal and valence tend to be considered enough for the so-called "core affect" (Russell, 2003).

In the circumplex model of arousal and valence, it is the combination of these two continuous dimensions that defines affective states (Figure 1.2). Therefore, pleasurable affective states correspond with high values of valence and low or high arousal, such as calm or excited, respectively. On the opposite end, unpleasant affective states correspond with low values of valence and low or high arousal, such as tired or tense, respectively. However, it is worth noting that the presumed orthogonality of these two dimensions has been called into question, as collections of ratings on these two dimensions show a relationship between both (Lang et al., 2008; Kurdi et al., 2017), as also confirmed experimentally, revealing a V-shaped relationship in which arousal is higher towards the extremes of valence, and lower towards neutral valence (Kuppens et al., 2013). Indeed, a study on brain activity using functional magnetic resonance imaging (fMRI) found that arousal is not separable from valence (at least, to predict arousal-related neural activity), suggesting arousal as either intensity of bipolar valence or as a linear combination of unipolar pleasant and unpleasant (Haj-Ali et al., 2020). Additional evidence from multiple neurological studies supports the circumplex model of affect (Posner et al., 2005).

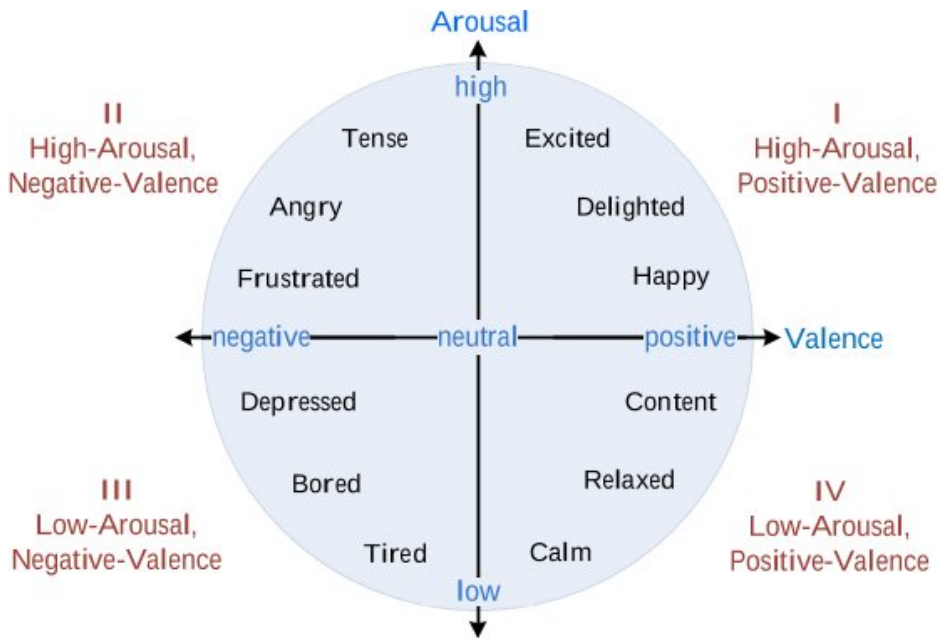


Figure 1.2: Circumplex model of affect

Graphical representation of the model proposed by Russell, as defined by two dimensions: arousal and valence. Adapted from Liu et al. (2018).

1.1.3 Neurophysiology of Emotion

The autonomic nervous system (ANS) is regulated by the brain to, in turn, unconsciously regulate bodily functions. The ANS is broadly considered crucial for emotions, including their generation, expression, experience, and recognition (Levenson, 2014). The ANS is typically divided into the sympathetic nervous system (SNS) and parasympathetic nervous system (PSNS) (McCorry, 2007). These two systems act as antagonists of each other, with the SNS promoting a “fight-or-flight” response corresponding to increased activation, and the PSNS promoting a “rest-and-digest” corresponding to calming (Porges et al., 1997; Levenson, 2014). As such, for example, the SNS increases heart rate while the PSNS decreases it, and the SNS dilates pupils while the PSNS constricts them (Kreibig, 2010;

Levenson, 2014).

Focusing on the brain, advances in neuroscience have identified a group of brain structures known as the limbic system that are particularly relevant for emotions. The limbic system is a network of structures in the brain that plays a key role in emotional processing and regulation (Rajmohan and Mohandas, 2007). Some of the more relevant structures identified as part of the limbic system include the prefrontal cortex, cingulate cortex, amygdala, hippocampus, nucleus accumbens, and hypothalamus (Morgane et al., 2005). Furthermore, the so-called emotional pathways have been identified, detailing the neural connections between different brain structures, both cortical and subcortical (Tamietto and de Gelder, 2010) (see Figure 1.3).

The ‘emotion system’ spans cortical and subcortical areas. Key subcortical areas include the amygdala (AMG) in the temporal lobe, the substantia innominata (SI) in the basal forebrain, and the nucleus accumbens (NA) in the basal ganglia, as well as nuclei in the brainstem including the periaqueductal gray (PAG) and the locus coeruleus (LC). Meanwhile, cortical areas include the orbitofrontal (OFC) and the anterior cingulate cortex (ACC). This system also has interconnections with visual pathways, such as the superior colliculus (SC) connecting to the amygdala via the pulvinar (Pulv) (Tamietto and de Gelder, 2010).

The amygdala is generally considered one of the most important brain areas involved in emotion processing. It has been identified as the key area for threat detection. In particular, the amygdala has historically been highlighted as a center for ‘fear’ processing. Many studies have established the key role of the amygdala in fear conditioning, creating memories of emotional, fearful events (LeDoux and Hofmann, 2018; Moscarello and LeDoux, 2013; Fanselow and Wassum, 2016). However, more recent research shows the amygdala’s involvement also in positive emotions (Xiu et al., 2014), establishing it as a more general structure of emotional processing. Therefore, current literature proposes two opposing groups of circuits in the amygdala: one for positive valence and another for negative valence (Beyeler et al., 2016; Namburi et al., 2015; Berridge, 2019).

Regarding emotional learning in the amygdala, its basolateral (BLA)




Figure 1.3: Emotional pathways in the human brain

Brain structures involved in emotions along with their neural connections. ACC: anterior cingulate cortex; AMG: amygdala; LC: locus coeruleus; NA: nucleus accumbens; OFC: orbitofrontal cortex; PAG: periaqueductal gray; Pulv: pulvinar; SC: superior colliculus; SI: substantia innominata. Adapted from Tamietto and de Gelder (2010).

and central (CeA) nuclei are established as the responsible units. In fear conditioning, signaling from the lateral thalamus (LT) to the lateral amygdala (LT) controls fear behavior by conveying the association between a conditioned stimulus (e.g., a tone) and an unconditioned stimulus (e.g., a foot shock) (Barys et al., 2020). A fast subcortical pathway specific for fear, and not positive emotions, has been shown to exist from the pulv-

inar nucleus of the thalamus to the amygdala via magnocellular axons, as threat detection requires short latency responses for survival (Méndez-Bértolo et al., 2016). Although fear is crucial for survival, it must be regulated within a homeostatic range, to avoid excess, as in anxiety disorders, or deficit, as in exaggerated risk-taking. To maintain fear balance, the insular cortex (InsCtx) of mice integrates bodily feedback (such as heart rate), in the form of predictive sensory and interoceptive signals (Klein et al., 2021).

The connectivity of the hippocampus (HPC) to the BLA plays a role in observational fear, a process to empathically experience another's fear (Terranova et al., 2022). Regarding empathy, studies with humans have identified two systems in the prefrontal cortex: one for emotional empathy, dependent on the inferior frontal gyrus, and one for cognitive empathy, dependent on the ventromedial prefrontal (Shamay-Tsoory et al., 2009).

Whereas the BLA and the CeA play a crucial role in emotional learning regardless of valence (positive or negative), a pathway from the BLA to the NA has been shown to play a key role in behaviors related to positive valence, particularly reward-seeking (Stuber et al., 2011). Within the amygdala, the circuit from the BLA to the CeA mediates appetitive behaviors (e.g., seeking food or water) (Kim et al., 2017). The CeA on its own has also been shown to modulate food consumption through a positive-valence mechanism (Douglass et al., 2017). In the CeA, the release of oxytocin, a hormone that plays a role in social bonding and other behaviors, modulates inhibitory circuits to suppress fear responses and decrease anxiety levels (mediated by astrocytes, non-neuronal brain cells that performs a variety of functions) (Wahis et al., 2021). As for social reward, an amygdala-to-hypothalamus circuit has also been identified, showing a role of the medial amygdala (MeA) in promoting the positive reinforcement of social interaction and the release of dopamine, a neurotransmitter involved in reward, in the NA (Hu et al., 2021).

Overall, as mentioned earlier, these findings tend to support the circumplex model of affect (Posner et al., 2005). Here, we have seen that a large body of evidence from studies in affective neuroscience and other re-

lated areas shows the existence of circuitry for valence, arousal, and their interactions at the neural level. As recently stated, different pathways exist to process the valence component of a stimulus, with the amygdala and the nucleus accumbens as some of the key brain areas, but with potentially distinct neuronal populations for positive or negative valence (Posner et al., 2005; Beyeler et al., 2016; Namburi et al., 2015; Berridge, 2019). Similarly, for arousal, the amygdala is also thought to be a crucial brain component, but through connectivity with different brain areas such as, importantly, the thalamus (Posner et al., 2005; Haj-Ali et al., 2020). However, as previously mentioned, different studies demonstrate that a clear relationship exists between both dimensions (Kuppens et al., 2013), including in the neural domain, at least in subjective self-reports (Haj-Ali et al., 2020). Taken together, these results highlight the complexity of emotions in the brain, with a large number of brain areas and circuits involved. Furthermore, they show that Russell's circumplex model of emotions, or a variant of it, is a good framework to study emotions, as it relates to measurable brain mechanisms.

1.1.4 Measuring Emotion

Due to the nature of emotions, which are complex and multidimensional, measuring them as they happen in an individual tends to be a difficult task. Over the last decades, researchers have used a wide variety of methods to scientifically measure emotions. Generally, these can be classified into physiological measures, behavioral measures, and self-reports (Mauss and Robinson, 2009).

Physiological Measures

Physiological measures rely on collecting data from bodily responses. In general, emotions are associated with a number of different physiological responses. Overall, these are based on measuring states from the autonomic nervous system, as it balances responses of activation and relaxation through the sympathetic and parasympathetic systems, respectively.

This, in turn, can serve as an estimation of affective states (Shu et al., 2018; Szwoch, 2015).

In recent years, researchers have used a number of different physiological measures to assess different physiological responses. Popular techniques involve measuring activity from the brain, the heart, the skin, and the eyes, through electrical or optical sensors (Dzedzickis et al., 2020).

Electroencephalography (EEG) is a technique to measure the electrical activity of the brain at the scalp level, non-invasively. To achieve this, a special cap equipped with multiple electrodes located on standard positions is used. Then, activity at each electrode is measured in the form of voltage, changing over time. Computational methods to estimate emotional states from EEG data are very diverse and no standards exist (Torres et al., 2020). Nevertheless, many approaches rely on the extraction of features at different frequency bands and brain locations, and the usage of different analysis methods, often advanced machine learning techniques, to classify discrete emotions or estimate arousal and valence (Torres et al., 2020; Shu et al., 2018; Dzedzickis et al., 2020; Mauss and Robinson, 2009). The asymmetry in activation between the two cortical hemispheres, especially in the frontal area, has been proposed as an indicator of valence, with greater left-sided activation for positive affective states (Coan and Allen, 2004; Zheng et al., 2015), although other authors argue that this is in fact related to approach versus avoidance motivation (Harmon-Jones, 2003; Mauss and Robinson, 2009; Bos, 2006). Meanwhile, alpha and beta are two frequency bands that are often used as indicators of arousal (Dzedzickis et al., 2020; Bos, 2006). Alpha activity generally corresponds to lower frequency brain waves and is associated with relaxation, while beta activity corresponds to higher frequency brain waves and is associated with activation or alertness. A reliable indicator of arousal often relies on the ratio of alpha and beta activity in the brain (Dzedzickis et al., 2020; Bos, 2006).

Neuroimaging techniques, such as fMRI, allow for much higher specificity of the activation of brain regions than EEG and therefore are better suited for scientific studies linking brain activity to emotions, although

at the cost of portability and temporal resolution (Mauss and Robinson, 2009; Brooke and Harrison, 2016). An overview of the insights on emotion allowed by neuroimaging was offered previously in section 1.1.3, identifying different brain structures and connections involved.

Heart activity is another popular physiological indicator of emotional states. It is often measured using electrocardiography (ECG), a technique to measure the electrical activity of the heart through the usage of multiple electrodes. For emotion estimation, a common approach relies on computing the heart rate variability (HRV), which is the time interval variation between heartbeats (Zhu et al., 2019; Acharya et al., 2006). Multiple features of interest can be extracted from HRV, including in the time domain, with the average interval between peaks or their standard deviation, or in the frequency domain, decomposing the signal between high- and low-frequency components (Guo et al., 2016). ECG and HRV analysis has been extensively used for emotion classification or estimation, once again using discrete or continuous emotion models (Mauss and Robinson, 2009; Dzedzickis et al., 2020; Shu et al., 2018; Agrafioti et al., 2012; Selvaraj et al., 2013), as well as to infer higher-level user states, such as cognitive workload and fatigue (Thayer and Lane, 2009; Mohanavelu et al., 2017; Segerstrom and Nes, 2007; Greene et al., 2016a). Generally, a higher heart rate, typically measured in beats per minute (BPM), is indicative of a higher arousal state (Shu et al., 2018), whereas a high HRV is indicative of good mental health and emotion regulation (Kemp and Quintana, 2013).

An alternative method to measure heart activity is by measuring the blood volume pulse (BVP), which corresponds to changes in blood volume produced by heartbeats. This can be measured using photoplethysmography (PPG), a technique to measure BVP by emitting light and measuring its reflection (Allen, 2007). Since it is a more indirect way of measuring heart activity than ECG, it may not be as reliable or comprehensive in some cases, but it can be measured more easily, with the usage of portable devices (Allen, 2007). Similarly to ECG and HRV, this technique is often used to assess emotional states (Cosoli et al., 2021; Menghini et al., 2019).

Another of the most commonly used physiological measures to estimate emotion is electrodermal activity (EDA) or galvanic skin response (GSR). This corresponds to the variation of the electrical properties of the skin in response to changes in the activity of sweat glands (Critchley, 2002). It is generally agreed that EDA is more related to arousal than to valence, with a positive correlation between EDA and arousal (Cosoli et al., 2021; Szwoch, 2015; Mauss and Robinson, 2009; Dzedzickis et al., 2020; Critchley, 2002).

Pupil dilation is a reliable indicator of emotion that can be measured using visual analysis. It is also influenced by light, so that is a factor that must be controlled when using pupil dilation as an indicator of emotion. Furthermore, some medications can also modulate pupil dilation. It is directly controlled by the autonomic nervous system, with the parasympathetic system handling constriction and the sympathetic system handling the dilation. As such, it has been used as an indicator of arousal (Bradley et al., 2008; Stanners et al., 1979).

Other physiological signals often used in emotion assessment are electromyography (EMG), respiration, and temperature (Mauss and Robinson, 2009; Dzedzickis et al., 2020; Shu et al., 2018; Szwoch, 2015).

Behavioral Measures

A classical approach for emotion recognition is the analysis of facial expressions (Fasel and Luetttin, 2003; Ko, 2018). As humans, facial expressions are one of the main ways in which we communicate emotions to one another. Certain facial configurations are usually called “emotional expressions” (Krumhuber et al., 2013) and are classically considered to be associated with specific emotional states (Ekman and Oster, 1979; Barrett et al., 2019), generally following the model of basic emotions explained earlier (Ekman, 1992). Therefore, over the last decades, there has been a great interest in the development of computational techniques for automatic facial expression analysis Fasel and Luetttin (2003); Ko (2018). Conventional approaches involve computer vision techniques to detect faces and their predefined landmarks (e.g., eyes, eyebrows, nose, mouth)

in images (either static or frames from a video), extract relevant features (e.g., movement of landmarks or distance between them), and perform a classification of the corresponding emotional expression (Ko, 2018). More recent approaches involve the usage of more advanced machine learning techniques: deep learning, with convolutional neural networks as the most popular method (Ko, 2018). However, these methods might be limited by the underlying emotional theory regarding facial expressions of emotions, which, as mentioned previously, has come under scrutiny for their reliability to convey stereotypical emotions (Barrett et al., 2019).

Going beyond the face, whole body gesturing, in what we call “body language”, can also convey emotional qualities, as an important form of non-verbal communication (Noroozi et al., 2021; Inderbitzin et al., 2011).

Speech is another of the main ways in which humans communicate with each other, including not only information but also emotions. Our voices convey emotion through subtle variations in their tone, and we rely on that to understand the full meaning of what has been said (Akçay and Oğuz, 2020; Koolagudi and Rao, 2012). Nowadays, computing systems that listen to our voices are mainstream, with so-called virtual assistants like Siri (Apple, Inc.) and Alexa (Amazon, Inc.). However, so far, they only recognize our words (i.e., what we say) and not our emotions (i.e., how we say it) (Ali, 2020; Schuller and Schuller, 2021). Still, the field of automatic speech emotion recognition has developed a variety of analysis methods over the last two decades, especially in recent years with the advances in deep learning (Schuller and Schuller, 2021). These methods are nowadays able to detect affect in voices to either classify it in basic emotions or estimate their valence and arousal (Akçay and Oğuz, 2020; Ali, 2020; Schuller and Schuller, 2021; Koolagudi and Rao, 2012).

On its own, language also encodes affect: humans express emotions in the choice of words when speaking or writing. Focused on text, so-called sentiment analysis aims at automatic emotion detection, by analyzing the words employed using natural language processing (NLP) techniques (Nandwani and Verma, 2021; Mäntylä et al., 2018). Like for other measures that we have seen, there are two general groups of methods for text sentiment analysis: those based on a dictionary or corpus, with ag-

gregated affective ratings, and those based on machine learning (Serrano-Guerrero et al., 2015). While some basic methods give a wide classification between positive, neutral, or negative emotions, more sophisticated algorithms offer an estimation of arousal and valence. These methods are being used in e-commerce to analyze the opinions of consumers, as well as in research on social media platforms (Nandwani and Verma, 2021; Mäntylä et al., 2018; Serrano-Guerrero et al., 2015). However, the reliability of this method is affected by the natural complexity of language, in which context and subtle cues might affect the meaning of the same words.

An alternative to these measures relies on analyzing subtle features resulting from the interaction of users with their computing devices. It is an inexpensive alternative or complement to other methods, as this does not rely on specific devices, allowing users to interact as they normally would, using a standard mouse, keyboard, or touchscreen. Examples of this are the analysis of keystroke dynamics (Epp et al., 2011), mouse movement patterns (Lali et al., 2014; Schaaff et al., 2012), or touchscreens (Yang and Qin, 2021).

Self-Reports

Finally, self-report measures involve the use of questionnaires to assess how individuals are feeling at a given point in time. Self-report measures are generally used for their ease of administration and low cost. However, they also have limitations, including the need to rationalize answers, which participants might have difficulties with, and hesitance to give answers that could be considered socially undesirable (Ciuk et al., 2015).

A well-established tool for emotion self-report is the Positive and Negative Affect Schedule (PANAS) (Watson et al., 1988). It consists of two 10-item scales, with each item rated on a 5-point scale. Each item is an emotional term (e.g., “excited”, “upset”), that people are asked to rate according to how they are currently feeling. It measures positive and negative affect as two separate dimensions, and therefore it is not based on Russell’s circumplex model of emotions (Watson et al., 1988; Russell,

1980).

A popular tool for the self-report of affect is the Self-Assessment Manikin (SAM), proposed in 1994 as a method to report pleasure, arousal, and dominance (Bradley and Lang, 1994). The SAM is a non-verbal, pictorial tool that uses 5 drawings per dimension to represent different values on a discrete scale. Pleasure is represented with a character ranging from a frowning face to a smiling one, while arousal has a character with an increasingly large explosion at its center, and dominance ranges from a small to a large character.

Another popular tool, proposed as a potential replacement for the SAM, is the Affective Slider (AS), which was introduced in 2016 as a digital tool to report arousal and pleasure on a continuous scale (Betella and Verschure, 2016). It consists of two separate sliders with emoticons on each side of both: unhappy/happy faces for pleasure, and sleepy/wide-awake faces for arousal. It was developed as an alternative to SAM, simplifying the assessment by relying on digital-first methods, while still providing highly reliable results (Betella and Verschure, 2016).

1.2 Human-Computer Interaction

Human-Computer Interaction (HCI) is focused on the way users and computing systems exchange information with each other within a defined context. This field researches ways of interaction through different means and has greatly evolved together with computers themselves due to the enormous technological advances that they have experienced since the first digital machines were developed mid-20th century.

The area of multidisciplinary research called Human-Computer Interaction was born between the late 1970s and the early 1980s with the creation and popularization of the personal computer (Jacko, 2012). However, human-computer interaction as an activity has existed since the first moments a person interacted with a computer back in the late 1940s with the first generation of computers, which were based on vacuum tubes.

The tendency in HCI has been towards more natural ways of inter-

action, from managing the individual connections of cables in the first computers to simply talking to them with some contemporary devices.

The interaction with the first general-purpose computer, the ENIAC (1945), was based on reprogramming it for each desired task by means of managing the individual connections of the different cables, as well as activating or deactivating different dials and switches. It was a time-consuming task that took days for the group of operators in charge. Over the next years, as the field of computing established itself, computer keyboards started to become increasingly common during the 1950s and 60s to more easily and quickly input information into the computer. It was during the 1970s and 80s that the first personal computers started appearing, marking the beginning of the ubiquity of computing devices in workplaces, homes, and other contexts. These computers started to introduce graphical user interfaces (GUI), allowing for richer interactions thanks also to the apparition of the computer mouse (Grudin, 2017). As computers were more widely available and used not only by professional and technical users, the way users interacted with them became more relevant. Thus, the research area of Human-Computer Interaction started to establish itself.

Over the last decades, computers continued getting smaller, more powerful, and easier to use. This brought them to a quickly increasing market of consumers. While familiar computers were being refined and improved, new interaction paradigms such as augmented reality (AR) and virtual reality (VR) were also being developed. These took longer to reach the mass market but are now becoming more widely available (Muñoz-Saavedra et al., 2020).

Nowadays, computing devices are ubiquitous and, in developed countries, people tend to have several. The development of the smartphone is of particular relevance, as it has brought a very capable computer to most people's pockets. In recent years, AR and VR have also reached an increasing number of consumers (Muñoz-Saavedra et al., 2020).

Focusing on these technologies, a reality-virtuality continuum can be established, ranging from a pure real environment to a fully virtual environment (Milgram et al., 1994). Augmented reality is close to the real

environment but it augments it by adding virtual elements, while virtual reality replaces the real environment as much as possible in our senses (see Figure 1.4).




Figure 1.4: Reality-Virtuality Continuum

Representation of the reality-virtuality continuum as defined by Paul Milgram. It ranges from the real environment to the virtual one. Adapted from Milgram et al. (1994).

Mixed reality covers the middle range of the reality-virtuality continuum, seamlessly merging the real and the virtual environment in a way that allows for full interactivity. As with AR and VR, mixed reality can be achieved by using a specific head-mounted display (HMD). However, it can also be achieved by creating an immersive environment with the necessary sensors and effectors (Bernardet et al., 2010).

In the last decades, the field of HCI has continued to expand, not only with the expansion of augmented, mixed, and virtual reality, but also with new form factors of interactive devices (e.g., smart speakers, smartwatches), more natural ways of explicitly interacting (e.g., gestures, voice), more contexts with advances computing devices (e.g., cars, sports), and many other aspects (Hasan and Yu, 2017). One of the overarching trends that has been developing, under different variations such as “symbiotic interaction” (Jacucci et al., 2014) or “human-engaged computing” (Ren et al., 2019) targets a closer integration between humans and computing devices, exploiting the capacities and capabilities of both.

To achieve this tight human-machine coupling, it is necessary to go beyond *explicit* interaction to also include *implicit* factors (Ju and Leifer, 2008). Explicit interaction refers to the conventional interaction process, in which the user performs an action on the device intentionally, con-

sciously, and purposefully, using an input device such as a mouse, a keyboard, or a touchscreen. Meanwhile, implicit interaction occurs without the explicit awareness of the user and includes aspects such as the context and internal states of the user (Serim and Jacucci, 2019; Ju and Leifer, 2008). Internal user states that could be of interest to consider include attention, intention, interest, and, importantly, emotions.

1.3 Affective Computing

Although some studies combining human emotions with computing devices existed before, the field of affective computing was formally established by scientist Rosalind Picard with the publication of her seminal 1995 report at the MIT Media Lab (Picard, 1995) (later expanded further in a book (Picard, 1997)). There, she defined affective computing as “computing that relates to, arises from, or influences emotion”.

The main idea behind the affective computing research area is that computers with affective capabilities would allow for a more natural human-computer interaction, a trend that, as seen in the previous section, has existed since the very first computers. Thus, we would communicate with a computing device in a way that more closely resembles how humans communicate with each other, moving even beyond what we consider natural user interfaces nowadays.

Affective computing is a multidisciplinary field, combining knowledge from diverse fields, such as psychology, computer science, physiology, cognitive science, and neuroscience (Arya et al., 2021). As such, advancements in those fields have the potential to have a significant contribution to affective computing. In turn, advancements in the field of affective computing also have the potential to provide insight into those domains.

Depending on the specific affective capabilities of a computer, different categories of affective computing can be identified (Picard, 1995), as seen in Table 1.1. Category I corresponds to most computers, which do not possess any affective capabilities, while category IV corresponds to

computers that could both perceive and express emotions. This fourth category would have the most complete emotion-oriented computers, being the goal of affective computing for many purposes. However, most recent advances in the field correspond to categories II (can express affect but not perceive it) and III (can perceive affect but not express it). Computers with the capability of perceiving the affective states of users would allow them to better interact with them, by adapting to the specific human needs and states.

Table 1.1: Categories of affective computing

Adapted from Picard (1995).

Computer	Cannot express affect	Can express affect
Cannot perceive affect	I	II
Can perceive affect	III	IV

For computers to recognize affect, a subset of the techniques for measuring emotions outlined previously in section 1.1.4 can be employed. Specifically, methods relying on physiological and behavioral measures tend to be used as automated emotion recognition techniques (Dzedzickis et al., 2020).

By now, the field of affective computing is well established, with more than two decades of history and constant advancements. Algorithms for emotion estimation using a wide variety of sources, keep improving, becoming more accurate and sophisticated (Wang et al., 2022). This progress has been applied in a plethora of fields (Calvo et al., 2015; Aranha et al., 2019), including education (Yadegaridehkordi et al., 2019), driving (Zepf et al., 2020), gaming (Robinson et al., 2020; Argasiński and Węgrzyn, 2019; Kotsia et al., 2013), art (Kostoulas et al., 2017; Wang and Chen, 2020), marketing (Caruelle et al., 2022), human-robot interaction (Gervasi et al., 2022), and healthcare (Yannakakis, 2018; Woodward et al., 2020; Greene et al., 2016b).

However, despite these remarkable advancements, a number of challenges and future possibilities still lay ahead. Beyond continuing the re-

finement of the algorithms employed with more sophisticated models or larger datasets (Wang et al., 2022), the underlying theory and considerations are crucial. For example, the pervasive use of the basic emotions model, still widely used in affective computing despite its known limitations, has been questioned, suggesting that more complex features that are not usually (but occasionally) considered also need to be computed, such as engagement, boredom, confusion, and frustration (D’Mello and Calvo, 2013). Furthermore, cognitive states, such as the ones just mentioned or cognitive workload, should also be considered in addition to affect in order to create even more advanced systems, capable of full implicit understanding.

To overcome some of these intrinsic limitations of the field of affective computing, other paradigms have been proposed in more recent years, such as “symbiotic interaction” (Jacucci et al., 2014) or “human-engaged computing” (Ren et al., 2019) as mentioned before, which integrate affective computing with other aspects (Jacucci et al., 2015).

1.4 Thesis Outline

Throughout this introduction, we have seen an overview of some of the topics of greater relevance for this thesis. First, we have seen the fundamental role that emotions play in human experience and action. We have examined how the concept of “emotion” has changed throughout history and how it is understood today by science. A key aspect is the need for a holistic vision of emotions that considers both the body and the mind together. The circumplex model of emotions, which decomposes affect into arousal and valence as bipolar dimensions, is supported by a large body of neuroscientific studies, which have researched the brain circuitry that modulates them. Then, we have seen an overview of the multitude of methods that can be used to measure emotion, grouped into physiological measures, behavioral measures, and self-reports.

Secondly, we have reviewed a general outlook of the field of human-computer interaction, from the beginning of digital computing to the latest

paradigms that are starting to be employed nowadays. We have seen the emergence of new paradigms that aim to make human-computer interactions more natural and intuitive, which has been the aim of this field since its inception. One of the challenges to achieve this is a better understanding of the users, not only when designing a system, but also when using it. Currently, most interactive systems are passive, reacting only to explicit input from users (e.g., clicking a button to perform a particular action). Unlike most humans, these systems lack emphatic capabilities to dynamically and implicitly infer how the user is feeling to be able to react to it. In order to overcome this challenge, different research fields have been established in the last two decades. Among these, we highlight affective computing.

Finally, we have explored the field of affective computing, which brings the two previous topics together to bring emotions into human-computer interaction. As we have seen, despite the great diversity and advancements achieved in this domain throughout the last two decades, a lot of challenges and potential remain. Although some of these challenges could be tackled with more sophisticated machine learning models and richer datasets to improve recognition accuracy against benchmarks (Wang et al., 2022), other fundamental issues require a revision of the underlying theories and assumptions. These include considering more internal states of the users: cognitive states, which are not emotions but still play a fundamental role in interaction (e.g., attention, cognitive load). Additionally, some of the research in affective computing still relies on theories of emotion that no longer are aligned with the latest psychological and neuroscientific views. An example of this is facial emotion recognition, one of the most popular methods in affective computing, which relies on a simplified vision of basic emotions that generally fails to consider the broad variation that exists (Barrett et al., 2019).

Considering all of this, this thesis proposes going beyond the constraints that exist in fields such as affective computing, thinking in terms of *empathic* systems. Such a system would not only rely on explicit interaction but also on implicitly understanding its users, estimating different affective and cognitive states. This would bring a series of benefits over

conventional, more passive systems, including an improved user experience, in which a system is able to understand how users are feeling (by inferring their internal states) and react accordingly. For example, if it detects that a user is frustrated, it could provide help; if it detects that a user is overwhelmed, it could automatically reduce the amount or complexity of the information being presented. This would enable a new generation of interactive devices, more effective in providing a personalized experience and helping users, thus enhancing their efficacy, efficiency, and subjective experience (e.g., enjoyment). Overall, this would mean a more natural and intelligent interaction. To achieve this overarching goal, this thesis has not only developed the necessary technological architecture and a variety of methods but also offered a series of insights into these states, including the concept of affect as a whole. Throughout the next chapters, we will present the work that we have completed toward this goal.

In Part I, over three chapters, we will present a sensing architecture capable of collecting a series of measures that are necessary for the implicit understanding of human cognitive and affective states. This will serve as the basis for an empathic system, which would need to collect this user data. Chapter 2 presents the general architecture as it was integrated into an immersive mixed-reality environment. This architecture follows a distributed, layered, and modular design, allowing for flexibility in its usage. Considering this, Chapter 3 presents a version of this architecture adapted to a virtual reality experience for the simulation of neurodiversity, aimed at creating awareness about sensory overstimulation. Then, Chapter 4 presents another version of this general architecture, expanded for its usage for stroke neurorehabilitation by also integrating an exoskeleton and a functional electrical stimulation (FES) system, as well as an expanded user model: a digital user twin.

In Part II, we will see a series of methods that we empirically tested to obtain meaningful indications of human internal states. Chapter 5 presents two studies aimed at inducing unconscious processing of emotional stimuli in an immersive environment, with the aim of gaining insights into the mental processes behind this. Additionally, this serves as an application of the sensing architecture presented in Chapter 2. Next, in

Chapter 6 we will present a method to infer affect from keystroke dynamics, in which we show that different features correlate independently with arousal and valence. Next, Chapter 7 presents a methodology that was employed at the height of the COVID-19 pandemic to assess the impact of the quarantine lockdown on the mental well-being of the population based on affective ratings. Here, we showed that a shift towards more negative ratings might be indicative of a deteriorated emotional state. Chapter 8 follows this with an extended study taking a look at additional implicit measures. Here, not only we confirmed our previous findings but we also extended them to identify the lockdown impacts on different demographic groups. Furthermore, we showed the potential of the different implicit features collected (mouse movements, keystroke dynamics, text sentiment analysis) to be indicative of these altered internal states. After this, Chapter 9 presents a novel methodology for collecting affective ratings through binary swiping on smartphones. We show its potential to collect rich affective information while also being faster and easier to use than previous tools. Here, we also capitalize on implicit interaction features by not only considering which way the users swipe but also how they do it: swipe dynamics (e.g., swipe velocity). Additionally, our results also offer insights into the circumplex model of emotions used throughout this thesis, joining existing research that challenges the orthogonality between arousal and valence.

Part III showcases two examples of interactive and adaptive systems that highlight the potential to enhance different kinds of systems through the implicit understanding of the user states. Chapter 10 is focused on an assistive system that was developed in the context of a law enforcement use case. It serves as an example of a simple system, capable of learning from successful interactions with users to autonomously provide suggestions to users when needed. Although this system is based only on online interaction data, we propose an extension with additional sensing devices to expand its capabilities. Chapter 11 presents a more advanced interactive and adaptive system, providing an experimental evaluation of its three core features: immersion, explicit interaction, and implicit interaction. Throughout three evaluation studies, our results demonstrate

that having an immersive environment and allowing users to directly interact is beneficial. Moreover, we show that having implicit interaction, inferring internal states based on physiological signals, to provide timely assistance enhanced the user experience quantitatively and qualitatively.

Finally, in Part IV, we will provide a general discussion of the insights obtained throughout this thesis (Chapter 12) and we will provide concluding remarks, including future work and the potential impact that this work can have on different fields, including the next generation of immersive interactive systems, affective science, and digital health (Chapter 13).

Part I

Architecture for Implicit Understanding of Internal States

Chapter 2

A DISTRIBUTED SENSING ARCHITECTURE FOR PSYCHOPHYSIOLOGICAL EXPERIMENTATION IN VIRTUAL AND MIXED REALITY

This chapter is based on:

López-Carral, H., Omedas, P., Zucca, R., and Verschure, P. F. (2023d). A distributed sensing architecture for psychophysiological experimentation in virtual reality. *Manuscript in preparation*

Virtual reality allows for the setup of ecologically valid environments in which subjects can act as they would normally while maintaining control over key variables such as the presented stimuli. Experimentation in such an environment provides the opportunity for recording multiple physiological signals, coming from different sensors, in order to infer

implicit user states such as emotional state or cognitive workload. To do this, we have integrated and validated a sensing architecture that allows for synchronization and online analysis of multiple signals coming from different devices, which might be distributed over a local network, depending on the experimental needs. The modular nature of this architecture allows for custom usage of the needed layers. We implemented this sensing architecture in the eXperience Induction Machine (XIM), an immersive space developed to conduct experiments in virtual and mixed reality. We have successfully employed this sensing architecture in different behavioral and psychophysiological experiments, with satisfactory accuracy, latency, and minimal data loss.

2.1 Introduction

In the last two decades, the advances in virtual reality (VR) techniques have led to their application in a wide range of scientific fields, including psychophysiological experimentation with humans (Diemer et al., 2015; Meyerbröker and Emmelkamp, 2010). One of the main benefits of this approach is that it allows for the setup of ecologically valid environments where users can act and behave in life-like conditions (Parsons, 2015). VR supports naturalistic and contextually rich scenarios along with a high degree of control over the key variables, such as the stimuli presented.

An important aspect is not only the use of more naturalistic stimuli but also the fact that the brain controls a body acting in a rich environment (Verschure, 2016). This idea first motivated the creation of the eXperience Induction Machine (XIM), a mixed reality interactive space for behavioral and psychophysiological experimentation under ecologically valid conditions (Bernardet et al., 2011; Betella et al., 2012).

Recent improvements in hardware and software allow to record psychophysiological signals in ambulatory conditions and make use of these signals to infer both implicit and explicit user states and actions in almost real time. However, a challenge is how to integrate signals from different sources, and being able not only to record them for offline analysis

but also analyze them online for real-time interaction. Here, we have enhanced the XIM as a general purpose infrastructure to support the analysis of physiological signals in a broad range of behavioral studies under naturalistic conditions by integrating a new Distributed Sensing Architecture (DSA).

We have validated the DSA by conducting different experiments, in which we used multiple physiological sensors, with software components distributed between different computers, and sending event markers with millisecond precision.

Our results show that the DSA that we have implemented in the XIM provides the necessary capabilities and performance. During our testing, it successfully combined signals coming from different sensors that we integrated, maintaining satisfactory accuracy, latency, and minimal data loss.

2.1.1 State of the Art

Various solutions have been developed for the online analysis of physiological signals. However, most of them are focused solely on the processing of EEG signals specifically for Brain-Computer Interface (BCI) research (Cowley et al., 2016), and generally do not take care of synchronizing signals coming from different sources, such as electrocardiogram (ECG) or electrodermal activity (EDA). Examples of such platforms are BCI2000 (Schalk et al., 2004), OpenViBE (Renard et al., 2010), and BCILAB (Kothe and Makeig, 2013). Another framework, SSI (Social Signal Interpretation) (Wagner et al., 2011), provides the tools for recording signals from different sensors (not limited to neurophysiology), processing them and detecting high-level features, such as gestures or emotional states. However, SSI does not support a distributed design and it is limited in extensibility for hardware and software compatibility. Finally, the MIDAS (Modular Integrated Distributed Analysis System) framework (Henelius and Tornaiainen, 2018) offers the building blocks for implementing a modular system for the analysis of arbitrary signals. It is based on

lab streaming layer (LSL ¹), a protocol for synchronized acquisition of different signals (time series) over a network.

For the new sensing architecture developed here, we needed support for different physiological signals such as ECG, EDA, EEG, body movements, and more, following a distributed design. For these reasons, we decided to base our new sensing architecture on the previously mentioned LSL protocol, as it offers networking capabilities that allow for a distributed design, while also providing signal synchronization and the possibility of online processing and centralized recording. This led us to integrate the MIDAS framework for the online processing of signals since it is based on LSL and also supports an extensible and distributed design.

2.2 Architecture

The implemented sensing architecture, DSA, is illustrated in Figure 2.1 and is composed of several layers: a number of sensors capturing users physiological data; synchronization and routing of the signals acquired by the sensors; online analysis, and storage of data for offline analysis. It integrates with applications such as the ones that run the visualization and interaction of the main experiment (e.g., Unity), with the purpose of sending markers for specific events and exchanging data (e.g., the application receives certain data for real-time interaction).

The DSA is modular, so some layers might be ignored if they are not needed for a specific use case, allowing for different configurations depending on the specificity of the experimental protocol. For example, one might choose to exclude the online analysis layer if no real-time processing is needed. On the contrary, one might choose not to use the recording layer if only real-time processing for interaction is needed.

¹<https://github.com/sccn/labstreaminglayer>

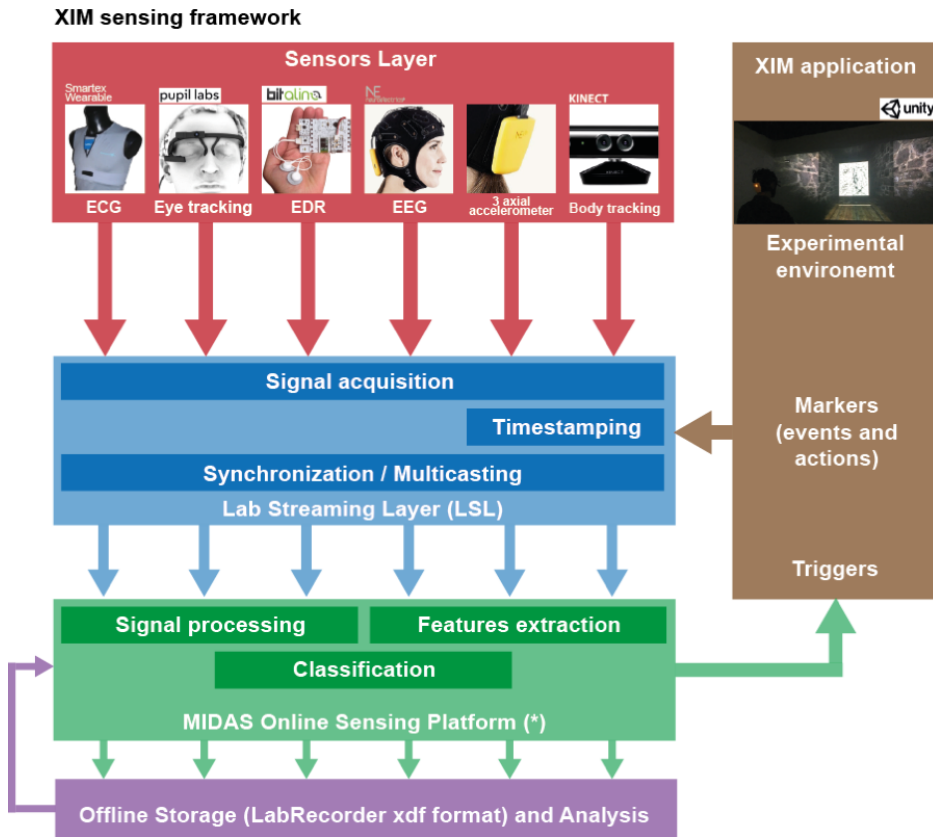


Figure 2.1: Implemented Distributed Sensing Architecture (DSA) as used in an experiment. It is composed of several modular layers: the signals from different sensors are acquired, synchronized, and streamed for on-line analysis and offline storage. External applications can receive these signals and send markers for events and actions.

2.2.1 Data Transmission and Synchronization

Data transmission and synchronization is based on the lab streaming layer (LSL) protocol ², originally developed at the Swartz Center for Computa-

²<https://github.com/sccn/labstreaminglayer>

tional Neuroscience (SCCN). LSL provides the necessary framework for streaming signals (times series) acquired from different sources within a local network, while also including specific metadata such as timestamps for synchronization. This process of data transmission and synchronization does not affect the quality of the signals and does not introduce artifacts. LSL is cross-platform and OS-independent and provides the necessary tools to develop plugins to use it with most sensors if those do not already exist.

LSL is based on an outlet and inlet transmission model. Outlets make streams of data available in the local network, while inlets connect to specific streams to retrieve their data. LSL streams are defined by parameters such as the name, the type, the sampling rate or the data type, which are set when creating the stream. These parameters are then used by the inlets to resolve the stream of interest.

Given this approach, based on UDP, LSL streams can be simultaneously received by multiple inlets in the local network, without needing to specify information such as the IP address of either the inlet or the outlet.

2.2.2 Sensing Devices

LSL can directly integrate the most common hardware via the plugins provided by most of the hardware's developers or through custom code.

Neurophysiology Sensor

For electroencephalography (EEG) data we used the Enobio headset (Neuroelectronics, Spain ³; 20 and 32 electrodes models). This device acquires 24-bit data at 500 Hz and is compatible with gel, dry and solid-gel electrodes. It has wireless connectivity using Bluetooth or Wi-Fi (depending on model). This offered out-of-the-box integration with LSL (in addition to TCP/IP). Along with the EEG data, it also provided 3-axis accelerometer data.

³<https://neuroelectronics.com>

In addition to observing brain activity for traditional EEG analysis, we can use this data to infer emotional states (Girardi et al., 2017; Soleymani et al., 2016; Bos, 2006).

Eye Tracker

We used the wearable Pupil eye tracker (Pupil Labs, Germany ⁴) to obtain gaze data and pupil dilation. This device offers binocular eye tracking using two adjustable cameras (a monocular model is also offered) with a sampling rate of 200 Hz, plus a world camera for egocentric vision and a sampling rate between 30 and 120 Hz depending on the image resolution. It connects to a computer via USB. It offers LSL integration using a plugin provided by Pupil Labs.

We can use the gaze data to accurately know whether the subject was looking at a specific location during the experimental procedure (e.g., the displayed stimulus), as well as detecting saccades. The pupil dilation measures can be used to infer internal states such as cognitive load (Beatty and Lucero-Wagoner, 2000; Pomplun and Sunkara, 2003) or arousal (Bradley et al., 2008).

Electrophysiology Board

For some physiological signals, we used the BITalino board (Plux, Portugal ⁵). This device has support for several kinds of signals, of which we used electrodermal activity (EDA) and electrocardiography (ECG). It connects wirelessly using Bluetooth. It requires custom LSL integration with the provided API. For this, we developed a small Python tool that connects to the device and streams the desired data.

We also tested the system with an e-Health Sensor Shield (Libelium, Spain ⁶). It supports multiple sensors, including EDA and ECG. It connects to a computer using USB. Like the previous device, it requires cus-

⁴<https://pupil-labs.com>

⁵<https://plux.info>

⁶<http://libelium.com>

tom LSL integration and we have developed a Python tool that connects to the device and streams the desired data.

Like other physiological responses, EDA has been repeatedly shown to be modulated by sympathetic nervous system activity, reflecting changes in emotional and cognitive states (Critchley, 2002; Lang et al., 1993). Similarly, ECG can also be used to classify emotional states (Agrafioti et al., 2012; Selvaraj et al., 2013). A combination of both EDA and ECG has also been used for inferring internal states (Betella et al., 2014c).

Sensing Shirt

We implemented custom LSL integration for the Smartex sensing shirt (Smartex, Italy ⁷) by developing a small sender program in the C# programming language, integrating the sensor library with the LSL library. It connects wirelessly using Bluetooth. It can provide data like electrocardiography (ECG), acceleration and respiration.

As described, ECG can be used to infer internal states. Respiration has also been linked to emotional feelings and it has been used to differentiate between different states (Philippot et al., 2002).

2.2.3 Online Analysis

In order to perform the online analysis of signals, the MIDAS (Modular Integrated Distributed Analysis System) framework (Henelius and Tornainen, 2018) was integrated, as an additional and optional layer of the DSA. MIDAS allows for a distributed structure based on nodes that use streaming signals transmitted using LSL.

The architecture of MIDAS follows a client-server model centralized in a dispatcher, which takes care of receiving requests from clients using a RESTful JSON API over HTTP.

A node needs to be created for every signal to be used, with secondary nodes taking care of fusing data from different primary nodes, for example. Thus, we might have one node for EDA and another for ECG, with

⁷<http://smartex.it>

a secondary node that receives signals from both to compute a measure such as arousal.

The implementation of the different nodes for the analysis of the physiological signals is user-dependent, as MIDAS does not provide signal-analysis algorithms. This kind of functionality must be implemented in Python either with custom algorithms or with existing third-party libraries, for both artifact correction and computation of higher-level features from the signals.

Overall, this layer gives the system the capability of having a real-time interaction loop using implicit signals (or internal states) inferred from the analysis of physiological measures in this layer.

2.2.4 Recording and Offline Analysis

The program LabRecorder is provided with LSL to facilitate the aggregated recording of the different streams of interest. It displays all of the streams found in the local network and records the ones selected by the user. LabRecorder uses the Extensible Data Format (XDF⁸) to store the recorded data in a single file per recording. The resulting file is stored locally, organized in folders according to the path specified in LabRecorder.

XDF, also developed at the Swartz Center for Computational Neuroscience (SCCN), is an open source format designed for physiological signals. Currently, the tool for importing XDF is available for Matlab and Python.

Once the data has been extracted from an XDF file, it can be analyzed in the usual ways with different analysis tools, without being limited in any way by the methods employed in the DSA. Since the DSA does not apply filtering or artifact removal to the signals (although it can be implemented in the online analysis layer), this preprocessing will be needed in most cases for the offline analysis.

⁸<https://github.com/sccn/xdf>




Figure 2.2: Schematic illustration of the eXperience Induction Machine (XIM), top view.

2.2.5 Immersive Environment

The eXperience Induction Machine (XIM) is a room for virtual and mixed reality interaction used for behavioral and psychophysiological experimentation (Bernardet et al., 2010). It is where the whole DSA was implemented. It is a space of 5.5 by 5.5 m equipped with a number of sensors and effectors, including a luminous interactive floor equipped with pressure sensors, cameras, microphones, a sonification system, and four projection screens that surround the space (see Figure 2.2). The combination of the pressure-sensitive floor and the cameras allow for tracking of the users in the space (Mathews et al., 2007). The projection screens and the loudspeakers are used to present the experimental application to the participants.

The experimental application runs on one of the XIM's computers and exchanges data with the rest of the DSA. For example, an application developed with the 3D game engine Unity (Unity Technologies, CA,

USA ⁹⁾ would integrate the LSL library to be able to send markers for specific events, and could also receive data from the online analysis layer by performing requests.

The XIM acts as the foundation for the implemented DSA, providing the necessary experimental environment. It provides the space for the real-time interactive applications in which this sensing architecture can be used to its full potential. However, even though the DSA is tuned for this space with the goal of conducting ecologically valid experimentation, it is flexible and modular enough to be used in different physical environments.

2.3 Application

The DSA has been tested and validated in a series of experimental studies. In these, the architecture has provided an appropriate performance and has shown the advantages of its modular and distributed structure.

Two of these studies, reported in more detail in Chapter 5, focused on the usage of subliminal stimulation in a navigation task while simultaneously recording several physiological signals (electrodermal activity, pupil dilation, EEG). The goal of this project was to study the behavioral and physiological effects of exposure to emotionally negative subliminal stimuli, as well as to validate the DSA. In these studies, carried out in the XIM, the participants navigated a virtual maze while being exposed to subliminal stimuli (either neutral or negative, the latter being a spider image) at the bifurcation points. The participants received feedback in the virtual experience depending on the subliminal stimuli that were presented and their navigational decision. The DSA was used in the following way:

⁹⁾<https://unity3d.com>

2.3.1 Sensing Devices and Signals

Three sensing devices were used, producing streams for five different types of data, connected to different computers, plus an additional stream for event markers:

- The neurophysiology sensor (Neuroelectrics Enobio32) produced a stream for EEG and another stream for triaxial accelerometer data. This device was connected wirelessly to the recording computer (via Bluetooth in the first study and via Wi-Fi in the second one).
- The electrophysiology board (BITalino on the first study; e-Health Sensor Shield on the second one) produced a stream for EDA. This device was connected to the recording computer (via Bluetooth in the first study and via USB in the second one).
- The eye tracker (Pupil Labs Pupil) produced a stream for gaze and another stream for pupil dilation. This device was connected via USB to an additional computer which connected to the local network via Wi-Fi.
- The experimental application produced a stream of markers for events, which were sent at irregular intervals ranging from seconds to minutes apart, depending on the phase of the experiment.

The streams had different sampling rates, ranging from 100 Hz, in the case of the accelerometer, to 1000 Hz, in the case of the EDA stream. The markers stream had an irregular sampling rate. This was not a problem for synchronization, given that the transmission framework that we used, LSL, was designed to handle these cases. The features of the sensors are reported in Table 2.1.

2.3.2 Experimental Environment

The experimental application was executed on a separate computer of the XIM that was connected to the local network using a wired connection.

Table 2.1: Signals captured and transmitted during an experiment using the DSA

Source	Signal	Sampling rate (Hz)
Neurophysiology sensor	EEG	500
	Accelerometer	100
Eye tracker	Gaze	200
	Pupillary response	200
Electrophysiology board	EDA	1000
Experiment application	Markers	Irregular, arbitrary

This computer presented the experiment to the user while sending markers for events to be recorded. In the first study, it used accelerometer data from the EEG headset to control the interaction. This interaction was achieved using the raw accelerometer data, processed directly by the experimental application in Unity, to detect head movements forward and backward (pitch rotation, establishing ad-hoc thresholds of acceleration for that axis). In the second study, two wireless controllers (Pyrus PY-1, China) were connected to the same computer via Bluetooth, sending key-press events directly to the operating system.

2.3.3 Online Analysis

MIDAS was used in a version of the first study with the purpose of computing the arousal responses of the participants to different stimuli. These arousal responses were computed using both EDA and EEG, on command from the experimental application, and the results were reported back to it. For this, arousal with EDA was computed by comparing the amplitude of the signal over two overlapping time windows of 5 and 10 seconds after the stimulus was displayed for 5 seconds; arousal with EEG was computed at the same time with the same overlapping time windows, but comparing the ratio of beta and alpha power (Bos, 2006). The MIDAS

architecture used for this consisted on a node for EDA, another node for EEG, and the dispatcher, with the client being the experimental application developed with Unity. The computer that executed the online analysis was connected to the local network using a wired connection.

2.3.4 Offline Storage and Analysis

All of the streams were recorded and stored offline for posterior analysis using LabRecorder. The result was a single XDF file per participant, which aggregated the data from all of the streams synchronized. The computer that executed the recording application was connected to the local network using a wired connection.

The offline analysis of the different physiological signals started with a preprocessing of the data to filter noise and other artifacts since this is a step that was not done before the signal recording. Using the event markers as references, we computed several measures, such as increases in phasic electrodermal activity, changes in pupil dilation, and brain activity in different areas and frequency bands, in response to the stimuli that were presented to the participants.

2.4 Performance

LSL has been validated to offer adequate measures of synchronization accuracy and latency. In these tests, LSL has been shown to achieve a synchronization accuracy of 1 ms or lower (Grivich, 2013; Ojeda et al., 2014) and a mean latency of 7.9 ms, considering the delays between hardware and software events (Grivich, 2013).

In our testing, the DSA has been fully validated after its successful performance in the experiments described in the previous section. During these experiments, all of the different layers performed as expected. Using LSL, the different signals coming from a variety of sensors were streamed over the distributed local network making them accessible to different computers in it. These signals were stored offline without data

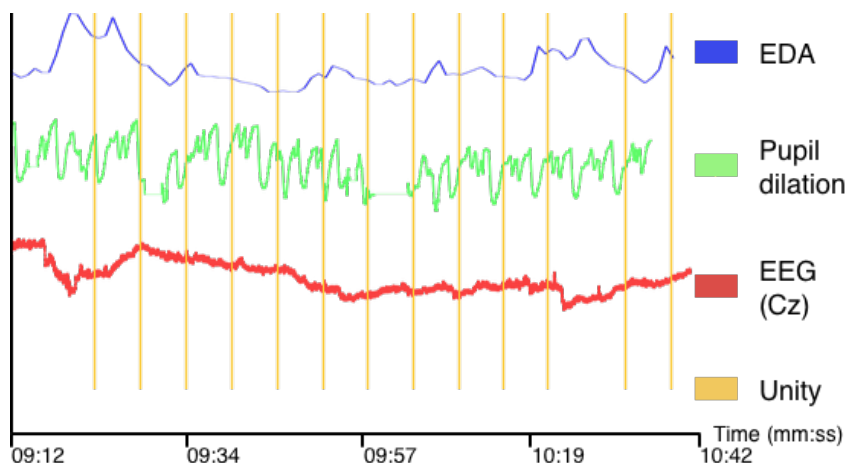


Figure 2.3: Plot of synchronized signals during an experiment. Three streams are shown: EDA, pupil dilation and EEG (for electrode Cz). Markers, corresponding to when subliminal stimuli (neutral or negative) were presented, are shown as vertical lines. The horizontal axis corresponds to the time since the beginning of the first trial for an example participant.

loss, accurately synchronized, into XDF files using LabRecorder for their posterior analysis. Furthermore, the MIDAS online processing layer performed analysis of high-level features from physiological signals, which it sent to the experimental application for its usage.

With an experimental setup like the one described in the previous section, running for around 30 minutes, a total of 1534 samples were sent at irregular rates, acting as markers. All of them were received, marking a data loss of 0%. These results were repeated for a total of 16 times without loss of information. This stream was synchronized with the simultaneous physiological recordings (see Figure 2.3). Moreover, the latency provided by the streaming system allowed for the signals to be used for real-time interaction. More examples of the successful performance of the DSA are provided throughout the following chapters.

2.5 Potential Applications

As shown, the DSA can be used as explained in psychophysiological experiments that require the recording and online analysis of multiple physiological signals in an environment that is more ecologically valid than traditional laboratory settings.

The online analysis layer of the DSA allows for a closed-loop human-computer interaction in which the physiological state of the users (as well as more complex inferred states such as arousal) directly affect their virtual environment.

An example of this would be brain-computer interaction (BCI). The online analysis layer (based on MIDAS) could be programmed to compute different features directly from a real-time EEG signal. These features could then be sent to a different computer on the local network to have an explicit effect on the user interaction, as research on BCI shows. One additional benefit of the DSA is the ease of including other types of signals in addition to just brain data, such as electrodermal activity, thus augmenting the information available for the interaction.

2.6 Conclusion

In this chapter, we have described a flexible Distributed Sensing Architecture (DSA) for psychophysiological experimentation in ecologically valid and immersive environments based on modern technologies. It allows for simultaneous online analysis and recording, with an organization based on different layers, some of which are optional to use. Since it is based on the lab streaming layer (LSL) system, the DSA can integrate signals coming from many different sources on a local network, which can also be received by a number of different machines for usage or recording. The online analysis framework implemented, MIDAS, also provides the possibility of distributing the different processing nodes within the local network. This online analysis offers the possibility of a closed-loop feedback using implicit signals inferred from the physiological measures.

We have shown that we implemented the DSA in the eXperience Induction Machine (XIM), an immersive space for mixed-reality, developed for its usage in behavioral and psychophysiological experimentation. Then, we used the DSA in the XIM to conduct some experiments, including one navigation task using subliminal stimulation while recording several physiological signals.

The results show that the DSA can successfully be used for real experimentation, thanks to high-accuracy synchronization, minimal data loss, and low latency. Furthermore, its implementation in a mixed and virtual reality space allows for experimentation in naturalistic settings not possible otherwise, thus increasing the ecological validity of the experiments. This will allow for new insights on brain mechanisms in more naturalistic settings than traditional laboratory experiments.

This has several implications both for psychophysiological experimentation and for virtual reality experiences. The successful implementation and usage of the DSA in this immersive environment highlights the potential of equipping users with diverse physiological sensors and other tracking devices. Such an augmented setup offers the possibility of tracking how users react during the VR experience in ways not possible otherwise. Beyond its usefulness in psychophysiological experimentation, the DSA offers the possibility of deeper human-computer interaction through the use of implicit features like user states inferred online with detailed contextual information (such as what the user was viewing or hearing, exact position in virtual space, etc.).

In the future, we expect to continue improving the system by integrating more physiological sensing devices and signals, as well as internal state estimation from explicit interaction, such as keystroke dynamics. Furthermore, we will continue with the deployment of the DSA in additional experiments. In Chapter 3, we present a version of this architecture employing different sensors in a virtual reality experience. Moreover, in Chapter 4 we present an extension of the DSA with additional devices and software components for its usage in stroke neurorehabilitation. Throughout Part II, we discuss different methods to estimate internal states based on different sources.

Chapter 3

A VIRTUAL REALITY SYSTEM FOR THE SIMULATION OF NEURODIVERSITY

This chapter is based on:

López-Carral, H., Blancas-Muñoz, M., Mura, A., Omedas, P., España-Cumellas, À., Martínez-Bueno, E., Milliken, N., Moore, P., Haque, L., Gilroy, S., and Verschure, P. F. M. J. (2022). A virtual reality system for the simulation of neurodiversity. In Yang, X.-S., Sherratt, S., Dey, N., and Joshi, A., editors, *Proceedings of Sixth International Congress on Information and Communication Technology*, pages 523–531, Singapore. Springer Singapore

Autism is a neurodevelopmental disorder characterized by deficits in social communication and repetitive patterns of behavior. Individuals affected by Autism Spectrum Disorder (ASD) may face overwhelming sensory hypersensitivities that hamper their everyday life. In order to promote awareness about neurodiversity among the neurotypical popula-

tion, we have developed an interactive virtual reality simulation to experience the sensory overstimulation that an individual with autism spectrum disorder may experience in a natural environment. In this experience, we project the user in a first-person perspective in a classroom where a teacher is presenting a lecture. As the user explores the classroom and attends the lecture, he/she is confronted with sensory distortions which are commonly experienced by persons with ASD. We provide the users with a virtual reality headset with motion tracking, two wireless controllers for interaction, and a wristband for physiological data acquisition to create a closed feedback loop. This wearable device measures blood volume pulse (BVP) and electrodermal activity (EDA), which we use to perform online estimations of the arousal levels of users as they respond to virtual stimuli. We use this information to modulate the intensity of auditory and visual stimuli simulating a vicious cycle in which increased arousal translates into increased sensory overstimulation. Here, we present the architecture and technical implementation of this system.

3.1 Introduction

3.1.1 Autism Spectrum Disorder

Autism Spectrum Disorder (ASD) is a neurodevelopmental disorder that affects social communication and is characterized by repetitive patterns of behavior. Individuals diagnosed with ASD may experience hypersensitivity, enhanced perception, and sensory overload (Mitchell and Ropar, 2004; Gomes et al., 2008). Some view this hypersensitivity as the result of hyperacute sensation; others, as a lack of prediction, leading to impairments in habituation. Regardless of the cause, these differences in sensory prediction, together with impairments in contextualizing sensory evidence, can handicap the understanding of others' actions and, consequentially, social interactions (Chambon et al., 2017).

3.1.2 Virtual Reality for Neurodiversity Simulation

With the goal of raising awareness among the neurotypical population about neurodiverse phenomenology, we developed an interactive virtual reality simulation to experience “neurodiversity”. In particular, we wanted to simulate the sensory overstimulation that people with ASD may experience during an ordinary situation. For the simulation environment, we have chosen a classroom given that it is a social context in which many possible stimuli may be present. In order to offer a realistic first-person experience, we chose to use virtual reality (VR) to place users in the perspective of a student affected by ASD (see Figure 3.1). Furthermore, we used a wearable device for acquiring physiological signals that we use to estimate arousal levels, which we use in real-time to create a closed feedback loop.



Figure 3.1: Screenshot of the classroom environment. The user is placed sitting at a desk, surrounded by other peers, and in front of a teacher, who gives a lecture on astronomy.

ASD encompasses a wide range of traits. As the use case of our project, we have simulated the experience of a teenager, focusing on Level 1 of the 5th Version of the Diagnostic and Statistical Manual of Mental Disorders (DSM-5) (Association et al., 2013) (“Requiring Support”); that is, although diagnosed with ASD, this person would not suffer severe deficits. The reasons to choose this level are because more advanced

levels would deal with more complex motor symptoms (Goldman et al., 2009), making it more difficult to simulate the experience; and because individuals with more severe symptoms (such as impaired intelligence or impaired communication) could even seek sensory stimuli.

Previous Examples The experience we are proposing is informed both by scientific literature and existing multimedia projects for ASD awareness. The available examples can be classified considering their format and level of interactivity.

- **Videos (regular):** In this type of experience, users can watch in a first-person view what a person with ASD would be experiencing. Examples (all but the first one are homemade): Carly's Café ¹, Walking Down the Street ², Sensory Overload Stimulation ³, Autism: Sensory Overload Stimulation ⁴.
- **360° videos:** In these ones, the viewer can also experience a 360° representation of their surroundings. Examples: Project Cape ⁵, The Party ⁶, Autism TMI Virtual Reality Experience ⁷.
- **Interactive:** This kind of experiences also allows users to interact with the environment. Examples: Auti-Sim (game) ⁸, Autism Reality Experience ⁹.

As mentioned, a variety of works have been created through different technological means to raise awareness about the sensory overstimulation

¹<https://youtu.be/KmDGvquzn2k>

²<https://youtu.be/plPNhooUUuc>

³<https://youtu.be/BPDTEuotHe0>

⁴<https://youtu.be/IcS2VUoe12M>

⁵<https://youtu.be/ZLyGuVTH8sA>

⁶<https://youtu.be/OtwOz1GVkDg>

⁷https://youtu.be/DgDR_gYk_a8

⁸<http://gamejolt.com/games/auti-sim/12761>

⁹<https://www.training2care.co.uk/autism-reality-experience.htm>

that someone with ASD could suffer. However, to the best of the authors' knowledge, none of the existing systems includes biofeedback to more realistically and dynamically recreate that experience. By using multiple physiological signals in real-time, we aim at overcoming this limitation and thus deliver a more complete simulation.

Affective State Estimation from Physiological Signals The autonomic nervous system modulates physiological responses, such as heart rate, respiration, and pupil dilation. This is directly reflective of certain internal human states, such as emotions and cognitive load. Thus, it is possible to use a variety of sensors to measure different physiological signals, such as the electrical activity of the heart using an electrocardiogram (ECG) or the skin's electrodermal activity (EDA), to learn about the users' states. In particular, these signals are known to correlate with affective states such as arousal (Szwoch, 2015), including in VR experiences (Betella et al., 2014b).

Electrodermal activity is the fluctuation of the electrical properties of the skin as modulated by sweat gland activity. This is controlled by the sympathetic nervous system in correlation with arousal (Critchley, 2002). Heart rate variability (HRV) is a measure of the variation of time intervals between heartbeats (Acharya et al., 2006), which can be derived from ECG data or photoplethysmography (PPG) data (Allen, 2007) and also correlates with arousal levels (Agrafioti et al., 2012). We use EDA and PPG together for increased robustness.

In this article, we present a novel virtual reality experience to simulate the sensory stimulation that neurodiverse people might face in their daily lives in order to promote awareness of this among the neurotypical population. This experience is enhanced by biofeedback using physiological signals to dynamically adapt the experience. Here, we describe the outcome, focusing on the stimuli used and the implementation in terms of its architecture and the estimation of users' internal states, before discussing the resulting work.

Table 3.1: Examples of stimuli used in the experience, divided between auditory and visual

Type	Stimulus	Examples
Audio	Background noise	Peers talking
Audio	Sudden noise	Car horn
Visual	Color	Shiny colors
Visual	Distortions	Moving patterns
Visual	Light	Excess of light

3.2 The Neurodiversity Experience

3.2.1 Stimuli

While immersed in the virtual reality experience, users are exposed to a series of stimuli whose properties (such as intensity and duration) are manipulated to induce a state of sensory overstimulation. The chosen stimuli are informed by a body of research on sensory overload in ASD and self-reports from individuals in the ASD spectrum. Considering the types of sensory overstimulation, the stimuli can be divided between visual and auditory (see Table 3.1).

Apart from being triggered, these stimuli can be modulated in intensity within a continuous range of values. Thus, they can be regulated depending on a number of factors, including the arousal levels of the users as inferred using their physiological responses.

3.2.2 Implementation

We have developed the “Neurodiversity Experience” as an interactive virtual reality experience augmented by biofeedback using a wearable device and implemented via a combination of different hardware and software technologies.

Architecture As a platform for the VR experience, we chose the Oculus Rift S headset (Oculus from Facebook Technologies, U.S.A.), a head-mounted display that provides the audiovisual experience to the users, as well as handling body movements (particularly head) and integrating two wireless controllers for interaction.

We engineered the virtual environment and the foundation of the experience using the Unity real-time development platform (Unity Technologies, U.S.A.). Using Unity, we developed the 3D environment in which users are situated during the experience to perceive a series of stimuli. This environment is populated by human-like characters, including other students and the teacher. They are animated realistically, in terms of both body movements and facial expressions. In the case of the teacher, the avatar moves around the classroom while gesturing, simulating the delivery of a lecture on astronomy. Mouth movements of this character are synchronized with a recording of the speech, performed by a human actress.

This 3D application is also the basis for the interaction process, taking care of integrating both explicit interaction, such as body movements and actions with the controllers, and implicit interaction, deriving mental and affective states from physiological signals.

The sensor used to acquire physiological signals is the Empatica E4 wristband (Empatica Inc., U.S.A.), a wearable device equipped with multiple sensors, including a photoplethysmography (PPG) sensor and an electrodermal activity (EDA) sensor. It offers the possibility of real-time data acquisition and streaming using wireless connectivity via Bluetooth to a computer (see Figure 3.2).

In order to process the physiological signals online and estimate the internal states of the users for interaction purposes, we developed an architecture integrating several software technologies (see Figure 3.3). We use the existing E4 streaming server ¹⁰ to forward real-time data using TCP socket connections. We developed a Python script that connects to that server, obtaining all data acquired by the wristband and relaying it

¹⁰<https://developer.empatica.com/windows-streaming-server.html>



Figure 3.2: Setup of the experience. The user is wearing an Oculus Rift S headset and an Empatica E4 wristband. The computer screen allows observers to see what the user sees.

using the lab streaming layer (LSL) system ¹¹, a protocol for streaming data which handles the networking and time-synchronization of signals for both online usage and recording. Then, we use the MIDAS (Modular Integrated Distributed Analysis System) (Henelius and Tornainen, 2018) to perform the online analysis of the signals streamed using LSL. To do this, we developed a node for each of the signals of interest (PPG and EDA), integrating the necessary analysis functions to estimate arousal levels. The virtual reality application then performs requests using a REST JSON API at regular intervals to obtain the processed arousal levels, which are used to modulate the intensity of the stimuli presented to the users.

Online Physiological Signal Analysis In order to estimate the arousal level of the users, we use a combination of two physiological signals: photoplethysmography (PPG) and electrodermal activity (EDA). From the blood volume pulse (BVP) measured by the PPG sensor, we derive heart rate variability (HRV). EDA and HRV are used in conjunction to estimate arousal levels.

¹¹<https://github.com/sccn/labstreaminglayer>

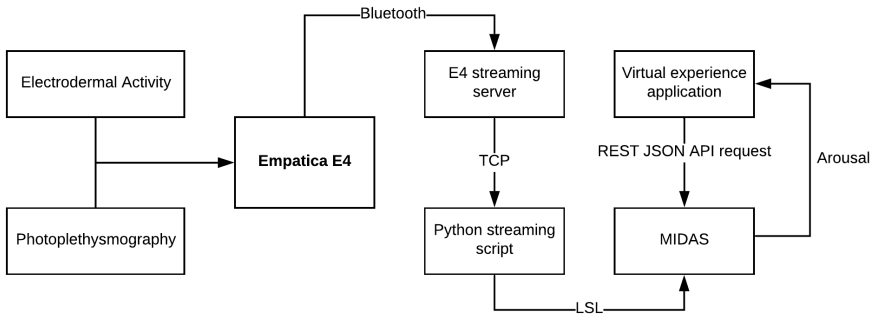


Figure 3.3: Architecture to process the physiological signals to estimate arousal levels online for interaction with the virtual reality experience. The physiological signals from the Empatica E4 wristband are transmitted to be streamed using LSL, to be analyzed online by MIDAS to infer arousal levels for a closed-loop interaction.

To compute the changes in the physiological signals, we use a moving average algorithm based on two overlapping time windows. The shorter time window, corresponding to the last 10 seconds, is compared to a longer time window of 30 seconds which includes the shorter window. By dividing the mean value during the short window over that of the long window, a measure of change is computed, centered around a value of 1. Values over 1 indicate an increase in arousal, while lower values denote a decrease.

This moving average is computed for each signal type individually, on the analysis node corresponding to each. Then, an additional node combines the result from the physiological processing nodes by computing an average that will act as the estimation of arousal levels (Betella et al., 2014a).

3.3 Discussion and Conclusions

We developed an innovative setup for a VR experience that places users in a classroom where they can assume the role of a student during a lesson. Throughout the experience, users are exposed to a series of stimuli to simulate their experience in ways that people with ASD may perceive them. To do this, we use a series of visual and auditory stimuli that are triggered depending on the timing and the actions of the users. The intensity of many of these effects is regulated using estimations of the arousal levels of the users, computed in real-time from physiological signals acquired by a wristband they are wearing, to further reinforce the experience using biofeedback for achieving increased effectiveness and realism. To accomplish this, we developed a software architecture that transmits the raw signals obtained by the wristband's sensors, processes them online, and makes them available for real-time usage by the VR environment to dynamically adapt it to the user.

The main implication of this experience is to raise awareness about the daily life of a student with ASD. To do so, this system will be deployed in several neurodiversity-related events, where users will be able to experience it. Moreover, it will allow us to understand the relationship between physiological signals, sensory overload, as well as attention and memory retrieval in classroom environments. This would be useful not only for gaining scientific knowledge and contributing to understanding the neurodiverse phenomenology but also for possibly helping teachers design more inclusive classrooms.

3.3.1 Further Steps

This article presents the technical implementation of our study focused on building an interactive VR experience targeting neurodiverse phenomenology. A further step in the validation of this experience will be to perform a user evaluation.

As a longer-term possibility, this experience could also support ASD individuals themselves. Previous studies have discussed the need for tech-

niques to improve predictive skills, rather than just treating ASD symptomatology. This could be done by adapting the type, intensity, and timing of the sensory overload stimuli to the degree of overload suffered by the individual.

Chapter 4

A SOCIALLY COOPERATIVE COGNITIVE ARCHITECTURE FOR REHABILITATION

This chapter is based on:

López-Carral, H., de la Torre Costa, J., Freire, I. T., Mura, A., and Verschure, P. F. (2023a). A socially cooperative cognitive architecture for rehabilitation. *Manuscript in preparation*

The economic burden of stroke, together with the COVID-19 crisis, is having a strong negative impact on healthcare delivery for those stroke survivors in need of continuous care and rehabilitation after hospital discharge. This situation poses a significant challenge for healthcare delivery and calls for rehabilitation interventions that can be deployed remotely. These interventions should be individualized to each patient's needs, following principles that promote motor and cognitive recovery. To achieve this, we have developed a socially cooperative cognitive architecture (SoCCA) capable of helping non-healthy users during Serious

Gaming rehabilitation routines by dynamically adapting aspects such as their intensity or duration based on the patient's state and performance. This cognitive architecture is based on a modular, distributed, and layered design and integrates signals from multiple sources, including different electrophysiological sensors, to infer the patients' internal states, such as arousal and fatigue. These variables, which form the dynamic user model (Digital Twin), are computed online based on the patients' health condition and physiological state. This approach could signify a significant improvement over traditional rehabilitation interventions, leading to better patient outcomes. Here, we describe the conceptual and technical design and implementation of this socially cooperative architecture.

4.1 Introduction

Stroke is currently the third cause of death and the main cause of adult disability worldwide (Kim et al., 2020). This implies massive health costs, especially under the current global health and financial crisis, which has resulted in the pressing need for remote neurorehabilitation interventions that can scale up to face this challenge (Wafa et al., 2020). Following a stroke, the dramatic loss of neural tissue leaves up to 70 % of patients experiencing persistent motor and cognitive impairments (Lai et al., 2002; Stevens et al., 2017). For instance, hemiparesis, or weakness in one entire side of the body, is one of the most common and disabling sequelae post-stroke.

Conventional rehabilitation approaches, such as occupational therapy, have focused on promoting the independence of patients when performing activities of daily living (Steultjens et al., 2003). However, these approaches often lead to compensatory strategies with a negative effect on neural remodeling and functional outcome (Jones, 2017). In addition, due to limited resources, these interventions may stop even when patients still have the potential to improve (Van De Port et al., 2006; Ballester et al., 2019). Indeed, it has been proposed that suboptimal recovery outcomes might lead to a phenomenon of deterioration of motor function

once discharged from hospital care. In this case, unsuccessful attempts to move the affected limb during recovery result in frustration and biased attention toward the unaffected, contralateral extremity. This phenomenon is widely known as *learned non-use* (Hidaka et al., 2012; Fuzaro et al., 2012; Taub et al., 1994). To prevent or reverse this behavior, Virtual Reality (VR) and other technologies that allow for remote interventions have emerged as recent treatment approaches in stroke rehabilitation, being rapidly adopted in clinical settings (Charles et al., 2020).

The use of Serious Gaming in the rehabilitation of motor impairments following cerebral damage has been intensely reviewed during the last years (Koutsiana et al., 2020; da Silva Cameirão et al., 2011; Maier et al., 2019). Importantly, recent studies that use specific VR interventions grounded on neuroscientific principles have shown that motor recovery is still possible in chronic stages (i.e., more than six months post-stroke onset) (Ballester et al., 2019), contrary to previous beliefs of improvement plateaus (Demain et al., 2006). This knowledge opens the doors for technology-based interventions that can be deployed in the homes of stroke chronic patients.

However, studies evaluating these interventions generally evaluate a one-size-fits-all strategy and usually do not distinguish between subgroups, treatment doses, or delivery modes. To overcome this, more recent therapy individualization approaches have focused on the so-called *Adaptive Difficulty* algorithms, which take user performance as the sole input. This strategy presents positive effects on the recovery and motivation level of participants (Pinto et al., 2018; Nirme et al., 2011; da Silva Cameirão et al., 2011) and is strongly grounded on the Yerkes-Dodson law, which poses an empirical relationship between stress and performance. Specifically, it says that performance increases with mental arousal up to a point, which might be considered the *optimal* level of challenge in learning tasks (Dodson, 1915; Ahmed, 2017).

However, by focusing only on the adaptation of the environment based on the patients' performance, we might be neglecting a whole set of potentially beneficial adaptations driven by recent research to deliver more effective rehabilitation interventions. For instance, the *embodiment* of the

virtual avatar might be promoted through its resemblance to the patients' anatomical features (Waltemate et al., 2018). Additionally, studies on the time-to-time variability of post-stroke fatigue suggest that the inference of patients' internal states may be crucial when designing personalized interventions (Lenaert et al., 2020).

Most of the novel approaches for personalized treatment rely heavily on Big Data management. An emerging technique that is being implemented in these growing data-driven healthcare practices is the so-called Digital Twin (Bruynseels et al., 2018). The notion of Digital Twin, with roots in engineering, describes the generation or collection of digital data representing a physical entity. It emphasizes the connection between this physical entity and the corresponding virtual counterpart, which is maintained by a real-time flow of information. In the context of healthcare, Digital Twin technologies have been applied to model the medical state of the patient, with the objective of facilitating diagnostics and helping in the design of more effective interventions (Croatti et al., 2020).

Here, we propose a socially cooperative cognitive architecture (SoCCA) to provide a principle-grounded and personalized rehabilitation intervention for stroke patients that can be deployed remotely. The SoCCA follows a modular and layered design to merge behavioral information from the patients' performance during a rehabilitation routine together with their physiological state and clinical profile. In particular, this system uses an existing cognitive architecture as its conceptual framework, hierarchically processing patients' information to construct a virtualization of them in the form of a Digital Twin. In turn, this information is used to optimize the rehabilitation sessions performed by the participants, dynamically adapting a series of Serious Gaming Scenarios, optionally completed with the help of an exoskeleton supported by functional electrical stimulation (FES).

4.1.1 Conceptual Framework

As its conceptual framework, this architecture follows the Distributed Adaptive Control (DAC) model (Verschure et al., 2003, 2014). DAC is

a cognitive architecture that provides a real-time model for perception, behavior, and cognition. It conceptualizes the brain as a control system that maintains a metastable equilibrium between the internal world of the body and the brain and the external world through action.

DAC proposes that goal-oriented actions emerge from the interplay of different processes that are organized in a four-layered control architecture with tight connectivity between and within layers, distinguishing: the Soma, Reactive, Adaptive, and Contextual layers. Across these layers, the architecture deals with the processing of states of the world, grounded in exteroception; the self, derived from interoception; and action, sensed through proprioception. The latter mediates between the first two via the environment.

The first level in the architecture is the Somatic Layer, which represents the body itself and defines the information acquired from sensation (from both internal and external stimuli), needs, and actuation (the control of the body's movement).

Subsequently, the Reactive Layer produces behaviors that support the basic functionality of the Somatic Layer in terms of reflexive behavior. The Reactive Layer constitutes the primary behavioral system based on the organism's physical needs and includes fast predefined sensorimotor loops (reflexes) that are triggered by low-complexity signals. In short, specific stimuli are hardwired with specific predefined actions.

Next, the Adaptive Layer extends the sensorimotor loops of the Reactive Layer with acquired sensor and action states associated with valence. The Adaptive Layer provides for the grounding of the representation of the world and the self through perceptual and behavioral learning systems. It acquires a state space of the agent-environment interaction combining perceptual and behavioral learning constrained by value functions to minimize perceptual and behavioral prediction error. Thus, adequate actions can be chosen to adapt to more complex inputs.

Finally, the Contextual Layer receives as its input the state-space acquired by the Adaptive Layer and generates goal-oriented behavioral plans and policies that can be expressed through actions. This layer includes mechanisms for short-term, long-term, and working memory, formatting

sequential representations of states of the environment and actions generated by the agent or its acquired sensorimotor contingencies in relation to the goals of the agent and its value functions.

The resulting architecture proposed here, following the DAC model, emerges from the interplay of the different layers, using multiple real-time signals to provide meaningful actions to assist patients during interactive rehabilitation activities (see Figure 4.1).

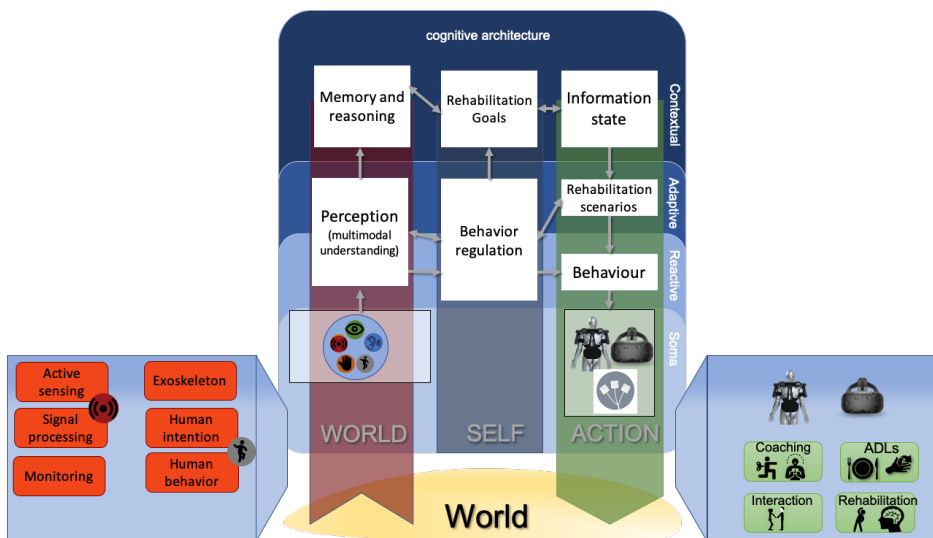


Figure 4.1: Diagram representing the general cognitive architecture based on the Distributed Adaptive Control (DAC) framework. This architecture is hierarchically organized into four layers: Soma, Reactive, Adaptive, and Contextual. A series of inputs are received, resulting in the perception of the *World*. This information is used to regulate the *Self*, which results in an *Action*. Applied here, the information from the patients is used to construct this architecture and result in assistive actions to aid in their rehabilitation.

4.2 Architecture

The socially cooperative cognitive architecture (SoCCA) proposed here uses a modular design to collect a wide range of signals, make them available for online processing to extract higher-order features, and then use these to update the Digital Twin and adapt the Serious Gaming Scenarios. In order to achieve this, the SoCCA is distributed in four distinct layers, following the pipeline of signal processing required for closed-loop interaction and the DAC model (see Figure 4.2).

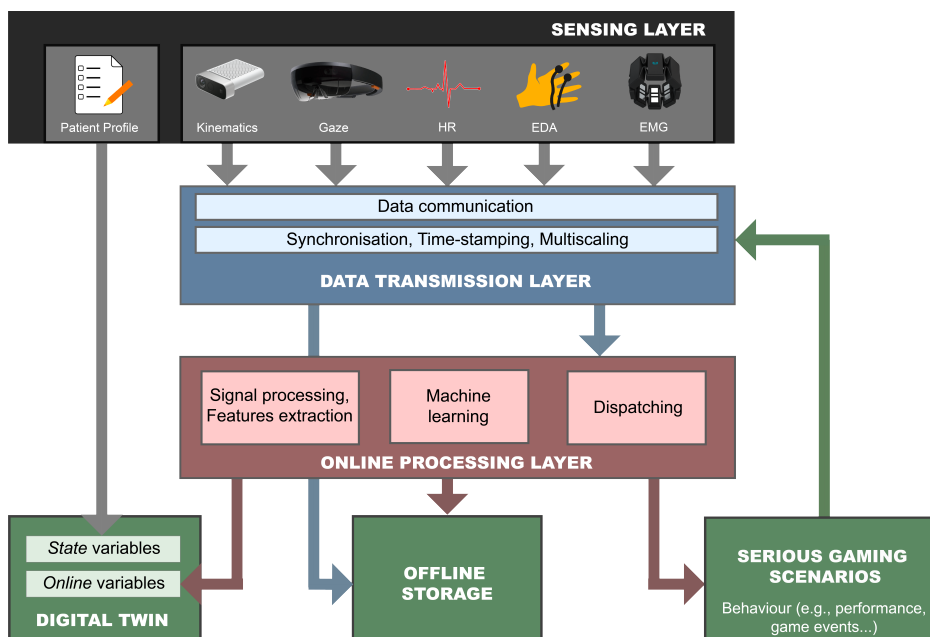


Figure 4.2: Schematic view of the socially cooperative cognitive architecture. It follows a modular design based on layers. The signals and other relevant information captured by the sensing layer are synchronized and transmitted to the online processing layer. This information is used to build the Digital Twin and dynamically adapt the Serious Gaming Scenarios.

4.2.1 Layered Structure

The first level is the sensing layer, which takes care of collecting information about the patient from various sources, thus acting as the Somatic Layer in the DAC framework. Some of the data may correspond to *state variables*, such as demographic or medical information, while some can be *online variables*, corresponding to real-time signals such as physiological measurements or body positions. Example signals include:

- **Kinematics:** information about the body movements of the patients, particularly the limbs used for interaction with the Serious Gaming Scenarios. This information is captured using a camera device with motion-sensing capabilities, such as the Azure Kinect (Microsoft, USA), and can be also collected from the exoskeleton used in this project. It is used to obtain information about the movement capabilities of the patient and to directly interact with the virtual environments through body movements.
- **Gaze:** information relative to where the patient is looking based on the position of their eyes, as measured by wearable eye-tracking cameras, such as the ones integrated into the mixed-reality headset HoloLens 2 (Microsoft, USA). This data can provide an indication of the patient's level of attention to certain aspects of the Serious Gaming Scenarios.
- **Heart rate:** physiological measure corresponding to the heart dynamics of the patient. This can be measured through different methods, such as measuring the electrical activity of the heart with electrocardiography (ECG) or changes in blood volume using light with photoplethysmography (PPG). These signals can be used to measure heart rate variability (HRV), which can be employed to infer the emotional states of the patients (Agrafioti et al., 2012; Dzedzickis et al., 2020). A possible device to measure this, which has been integrated into the SoCCA, is the Empatica E4 wristband (Empatica, USA).

- **Electrodermal activity (EDA)**: physiological measure about the changes in electrical conductance of the skin, which is indicative of psychophysiological processes regulated by the autonomic sympathetic system (Critchley, 2002). This signal can provide information related to emotional and cognitive states, including arousal (Wang et al., 2018) and cognitive load (Setz et al., 2009). This can also be measured using the Empatica E4 wristband.
- **Electromyography (EMG)**: technique to measure muscle movement based on electrical activity. Beyond providing information for gesture recognition (Zhang et al., 2011), EMG can also be used to estimate levels of fatigue (Cifrek et al., 2009). One potential device to obtain this information, also integrated into the SoCCA, is the Myo armband (Thalmic Labs, USA).

The second layer of the SoCCA corresponds to the synchronization and communication of the different signals obtained by the sensing layer to make them available for the rest of the components of the architecture. To achieve this, signals are transmitted using the lab streaming layer (LSL¹) system. LSL allows for the streaming of signals, handling the networking and their unified collection with time-synchronization, as well as online processing. Furthermore, LSL streams can be transmitted across a network and accessed from different applications simultaneously. Therefore, multiple computing devices might be used, allowing for a distributed architecture. This layer, allowing some simple responses to raw inputs, is analogous to DAC's Reactive Layer.

The third layer corresponds to the online processing of the signals transmitted by the previous layer. This online processing is performed with the aim of extracting higher-order features from the raw signals initially captured, including sensor fusion. An example of this is inferring the levels of arousal of a patient by combining the HRV and EDA signals. This online analysis is performed using the MIDAS (Modular Integrated Distributed Analysis System) framework (Henelius and Torn-

¹<https://github.com/sccn/labstreaminglayer>

ainen, 2018). MIDAS captures the desired LSL streams in specific processing nodes for each type of signal. Additionally, secondary nodes can be used for signal fusion and machine learning. Access to the nodes is handled by a centralized dispatcher, which manages the routing of the requests performed (using a REST JSON API) by client applications. This information is used by the final components of the SoCCA: the Digital Twin and the Serious Gaming Scenarios. This layer is equivalent to the Adaptive Layer defined by the DAC model.

4.2.2 Digital Twin

The goal of the SoCCA's Digital Twin is to generate, filter, and collect all the relevant information regarding each user of the system and integrate it into one single data structure. This information is used to flexibly modulate the human-robot interaction when an adaptable exoskeleton is used to assist the patients' movements, as well as the rehabilitation tasks to each patient's capacities and preferences over time. In other words, the Digital Twin creates a virtualization of the patient, which uses this information with the goal of optimizing a set of parameters both from the virtual rehabilitation scenarios and the exoskeleton interface.

The Digital Twin also follows DAC's hierarchical layered design, as it is composed of its own Reactive and Adaptive layers (see Figure 4.3). The Reactive Layer serves as an integration layer, gathering information from several input sources, whereas the Adaptive Layer processes these raw data in order to produce personalized suggestions to other modules of the SoCCA architecture.

More concretely, the input information gathered by the Digital Twin's Reactive Layer can be divided into state variables and online variables:

- **State variables:** define the profile of each user based on demographic and medical information. This information will be acquired from the end-users (patients) and clinicians via questionnaires and medical reports. This type of data is mostly static, as it will not vary throughout the therapy session (e.g., age, gender, preferences, baseline level of impairment).

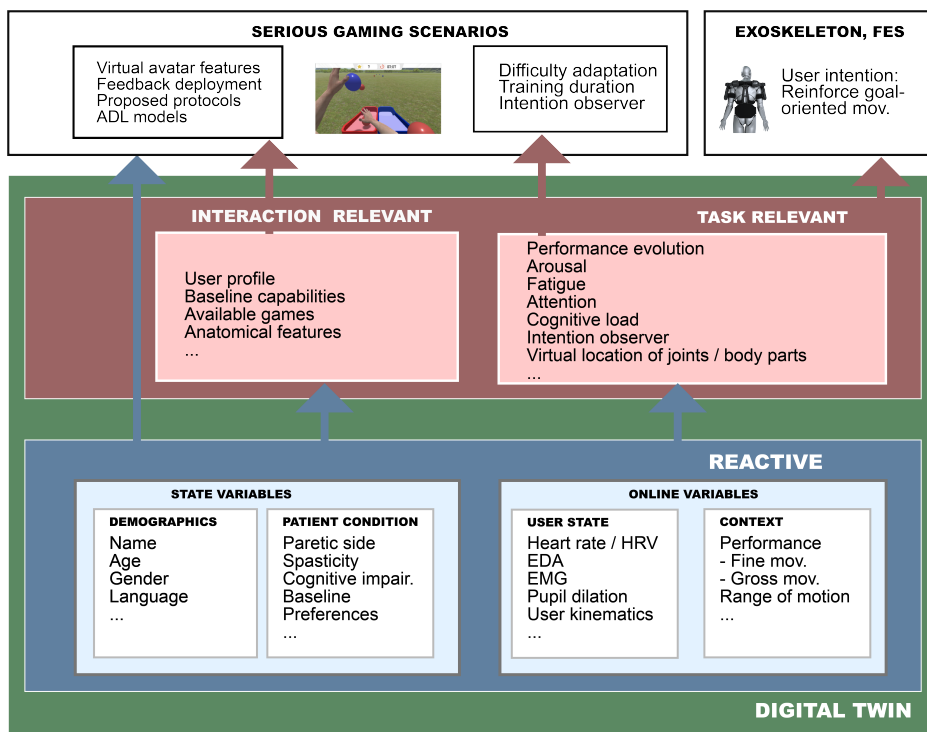


Figure 4.3: Digital Twin architecture. It comprises two levels of processing, divided into a Reactive Layer, handling online and state variables, and an Adaptive Layer, which provides actions that are relevant for interaction with the Serious Gaming Scenarios and for the tasks that the patient is carrying out, including adapting the usage of the exoskeleton through Functional Electrical Stimulation (FES).

- **Online variables:** comprise all the relevant user data that is dynamically updated in real-time over the course of the therapy session. They integrate information about the behavior (e.g., performance) and the internal states, and it is inferred through a multimodal probabilistic estimation from the sensors' data (e.g., arousal, fatigue, cognitive load).

The information gathered by the Digital Twin's Reactive Layer is used by its Adaptive Layer to generate interface- and task-relevant information. On the one hand, the interaction-relevant information provides the first layer of adaptation between the patient and the gaming system. It is generated from the state variables of the patient and allows the SoCCA to personalize general aspects of the gaming interface and the scenarios, from the aspect of the avatar to the type of feedback the patient will get (e.g., visual or auditory feedback, language). Based on the clinical profile of the patient, the Digital Twin will also adapt the available rehabilitation routines in order to promote those games that are more suitable to promote recovery. On the other hand, the task-relevant information produced by the Digital Twin provides a more coarse-grained level of adaptation to the patient that focuses on specific aspects of each task. The task-relevant information is computed from the online variables that the patient generates through real-time interaction with the system. Aspects such as the patient's performance evolution and internal variables like arousal and fatigue are computed in this layer. This information is used to dynamically modulate parameters of the game like the adaptive difficulty and the training duration. This layer also comprises the extraction of the motor kinematics and the generation of movement intention parameters that will be used by the SoCCA to modulate the behavior of the robotic exoskeleton.

The implementation of the Digital Twin is based on two main components: the Digital Twin API and the Digital Twin database. The database is the memory component of the Digital Twin architecture. Its function is to store all the information related to each patient and to keep it up to date. It is implemented as a document-oriented database using MongoDB ², where each Digital Twin profile is stored as a unique document. Each digital twin entry is initialized with the state variables acquired from the patient profile and questionnaires. Additionally, it also stores the main statistics of each interaction between the patient and the SoCCA.

Both the update of the database and the interaction with it are centrally controlled by the Digital Twin's API. The API's function is twofold: it performs the basic CRUD (create, read, update, and delete) operations

²<https://www.mongodb.com>

that keep the database up to date, and it is in charge of filtering the online and state variables to produce the task- and interface-relevant outputs of the Digital Twin. The API is written in the Python programming language and communicates with the database using BSON as the data interchange format.

4.2.3 Serious Gaming Scenarios

The SoCCA interacts with a set of VR Serious Games tailored for the motor and cognitive recovery of stroke patients, based on the Rehabilitation Gaming System (RGS). These Serious Games are grounded on a set of neurorehabilitation principles that target experience-dependent cortical plasticity mechanisms, shown to promote learning and recovery after brain injury (Maier et al., 2019). The adaptations driven by the architecture are specific to the protocol at hand. However, some common guidelines can be followed and generalized to each intervention.

The architecture takes *state variables* information from the *User Profile* to adapt, for instance, the anatomical features of the virtual avatar accordingly (i.e., color, size of the arms, gender). This does not only fulfill the embodiment priors of anatomical congruence (Ehrsson, 2012) but also helps the mapping between the real and computer-generated effector to be congruent, promoting agency and body ownership (Limanowski and Blankenburg, 2016). This information might also be used to adapt the feedback provided to patients with aphasia, characterized by impairments in the comprehension and/or expression of language (Vallila-Rohter and Kiran, 2013).

Online Variables are used to modify properties of the virtual environment dynamically. For instance, information from the performance can be used to modify the different variables that describe the difficulty of the Serious Gaming Scenario, following a similar approach as in (Nirme et al., 2011). In the case of the game presented in Figure 4.3, a set of colored spheres approach the patient, who must grab them and place them in the matching colored baskets. The difficulty of this game is modified through parameters such as *size*, *velocity*, and *dispersion* of the spheres,

or *number of colors* and *movement* of the baskets.

Internal states of the patients, such as cognitive or emotional conditions, are inferred by using different inputs based on existing knowledge in fields such as neuroscience, cognitive science, and psychophysiology. In order to estimate these states at a higher level of abstraction, the signals captured and transmitted by the initial layers of the SoCCA (such as heart rate, electrodermal activity, or muscle movements) are analyzed, either separately or combined. The result of this is a series of variables that form the Digital Twin by providing a representation of the patient's internal state, including arousal, fatigue, and cognitive load. This information is dynamically used to modify aspects such as the duration and intensity of the intervention. For instance, if the system detects an excessive level of fatigue, the Serious Game could be adapted to reduce the intensity of the ongoing task or even to finish the training session. Thus, the Serious Gaming Scenarios are capable of providing a high level of adaptability to each patient, both throughout the rehabilitation process and during an individual session.

4.3 Discussion

In this chapter, we have presented a novel Socially Cooperative Cognitive Architecture (SoCCA) capable of dynamically providing adaptations to Serious Gaming Scenarios for neurorehabilitation after a stroke. Currently, stroke is one of the leading causes of death and disability. It has been widely established that motor recovery reaches a plateau three to six months post-stroke (Demain et al., 2006), which has commonly justified the discharge of patients from hospital care at chronic stages. However, recent evidence using VR-based therapy shows that recovery is still possible beyond one year after the stroke onset (Ballester et al., 2019). It is believed that, after a stroke, shock leads to a learning process in which the brain progressively suppresses the use of the affected extremity (i.e., *learned non-use*) (Wolf et al., 1989). Moreover, stroke patients are prone to declines in mobility during chronic stages, likely due to inactivity and

comorbidities such as fatigue or depression (Van De Port et al., 2006). The *Stress and Coping* model (Folkman and Lazarus, 1988) suggests that *suboptimal* outcomes could, in turn, worsen due to self-limiting cognitive beliefs, leading to a cycle where stress, coping strategies, and function degrade recursively (Ballester et al., 2016). Thus, personalized remote interventions not only could promote activity during chronic stages but also an adaptive level of challenge and tailored manipulations in the interaction with the environment could result in reduced levels of stress and increased self-efficacy (Jones, 2006; Timmermans et al., 2009). Importantly, stroke patients exhibit a pronounced sensitivity to success and failure, which biases arm use (Ballester et al., 2016) and reemphasizes the need to adapt the rehabilitation protocols in an optimal way.

A recent report proposed a set of Universal Design recommendations specific to stroke survivors that include: “Make it fun, do not make people fail, empower and encourage. The technology needs to be highly adaptable to different sets of abilities [...]. Importance of seeing your progress, allow one-sided use (hemiplegia), avoid sensory and activity overload (fatigue), complement speech with images (aphasia), limit demand on memory, support learning and avoid errors (memory problems), and include multiple modalities in the design (reduced vision or hearing).” (Magnusson et al., 2018). We believe that the proposed architecture will be helpful in the design of future computer-generated stroke therapies that take these concepts into consideration. Additionally, thanks to the SoCCA’s modular and distributed design, it is possible to adapt its integration to each specific use case by, for instance, only using the desired components or extending it to use additional sensors or infer additional internal states.

The SoCCA is grounded on the Distributed Adaptive Control (DAC) theory as its conceptual framework, establishing a hierarchical layered structure. The first layer of the SoCCA captures patient data, making it available for its synchronized transmission, followed by its online analysis to infer higher-order internal states (e.g., arousal, fatigue, cognitive load). This information is then used to update the Digital Twin, corresponding to a virtualization of the patient that is then used to adapt the Serious Gaming Scenarios.

The integration of this architecture in neurorehabilitation routines allows for a more personalized treatment not possible otherwise. This could have a significant impact on the recovery outcomes of patients, going beyond what conventional, more static approaches allow. Furthermore, the use of the system proposed here, which can function autonomously without the constant supervision of specialized clinical personnel, could allow rehabilitation interventions to be deployed at home, thus increasing the number of patients reached and rehabilitation sessions performed. This is of special interest considering the impact of the COVID-19 outbreak, which negatively affected stroke care, resulting in a significant drop in admissions, thrombolysis, and thrombectomy. Patients not going to the hospital was one of the main limiting factors of this (Zhao et al., 2020).

Future work with the SoCCA includes its deployment in rehabilitation settings to perform user validations with stroke patients in clinical studies. This will involve the collection of data from each aspect of the architecture for online and offline analysis within the *Horizon 2020 Re-Hyb project*³. The results of using the system proposed here, in terms of rehabilitation outcomes, will be compared with those from classical rehabilitation routines, expecting a significant improvement in both functional outcomes and user satisfaction that highlights the benefits of the approach proposed here.

³<https://rehyb.eu>

Part II

Methods for Implicit Understanding of Internal States

Chapter 5

INDUCTION OF UNCONSCIOUS PROCESSING OF EMOTIONAL STIMULI IN AN IMMERSIVE SYSTEM

This chapter is based on:

López-Carral, H., Zucca, R., Omedas, P., and Verschure, P. F. (2023e). Induction of unconscious processing of emotional stimuli in an immersive system. *Manuscript in preparation*

While traditional human-computer interaction has usually relied on users' explicit input, novel approaches are exploiting their unconscious processes through psychophysiological measurements. Furthermore, standard decision-making protocols based on subliminal stimuli have been carried out under well-controlled laboratory condition. It is unclear how these results generalize under real-world conditions. Here, we investigate the effects of the presentation of subliminal emotional stimuli in an immersive virtual-reality environment both at the behavioral and physiological level. This approach allows us to overcome the limitations of

traditional laboratory settings by providing more ecologically-valid conditions. Participants were asked to navigate a virtual maze, where they encountered subliminal cues at bifurcation points. Using a novel framework for psychophysiological signal integration and real-time control of VR content, we recorded signals from a range of sensors worn by the participants including electroencephalography (EEG) for assessing the neural effects of these stimuli, while electrodermal activity (EDA) was used to derive arousal responses. Although the behavioral performance of the users in the navigation task was not consistently affected by the subliminal stimuli above chance, significant electrophysiological responses were detected. EDA increased significantly more when a negative affective stimulus was presented, as compared to a neutral stimulus. Similarly, we observed that an increased beta/alpha power ratio in the frontal electrodes is a good predictor of the level of the arousal induced by subliminal primes. Furthermore, the presentation of the negative subliminal stimulus produced a significant increase in spectral power within the high delta band in the fronto-parietal EEG electrodes. These results highlight the possibilities of tracking neurophysiological responses produced by adapting subliminal stimulation paradigms into ecologically-valid, virtual-reality contexts. We demonstrate that it is possible to analyze human brain dynamics during active navigation in enriched environments while, at the same time, gaining insights about the processes underlying conscious and unconscious perception, action and cognition.

5.1 Introduction

5.1.1 Implicit User States in Human-Computer Interaction

Traditional Human-Computer Interaction (HCI) relies on the users' explicit input via keyboards, pointing devices, and many other tools. However, in recent years, novel approaches taking into account the users' unconscious states have been presented (Andre, 2013; Wagner et al., 2013).

The goal of developing a system that uses a variety of implicit user's inputs would be a symbiotic system: a system in which human cognitive capabilities are augmented by complementing them with those of the machine. Ideally, this would result in a system that improves the interaction in ways not possible otherwise.

To be able to develop such class of empathetic systems it is necessary to understand the ways in which a computing device can communicate with the users using these unconscious processes. In particular, it is necessary to understand how to appropriately interpret and affect the user's internal states.

To address this issue, different sensors can be used to record physiological signals and infer implicit user states. Electrodermal activity (EDA) has been used to measure arousal levels (Betella et al., 2014c; Carbonaro et al., 2012; Courtney et al., 2010; Picard et al., 2001). Electroencephalography (EEG) has also been used to detect emotions (Bos, 2006; Kim et al., 2013), including during the presentation of subliminal visual stimulation (Bernat et al., 2001).

5.1.2 Unconscious Stimuli Processing

Decades of experimental research have shown that the human brain is capable of processing stimuli that are presented below a threshold of conscious perception, and thus influence behavior (Bornstein and Pittman, 1992). The response priming paradigm has been frequently used together with subliminal stimuli to unconsciously bias responses of participants (Schmidt et al., 2011; Eimer and Schlaghecken, 2003). Thus, when a participant is exposed to a stimulus for a few milliseconds, without conscious awareness, followed shortly by a target supraliminal stimulus, a stimulus-response association is developed (Damian, 2001), further reinforced by providing feedback (Hommel, 2000).

Another habitual component in visual subliminal priming is masking, used to disrupt the retinal afterimage right after the presentation of the subliminal stimulus (Vorberg et al., 2003; Kouider and Dehaene, 2007). This is usually some visual noise to control strictly the time that the sub-

liminal prime is perceptually available (Scharlau et al., 2006).

Besides behavioral biases or subjective reports, which can diverge from actual perceptual thresholds (Cheesman and Merikle, 1984), the effects of subliminal stimulation could be observed on physiological responses such as electrodermal activity (Lazarus and McCleary, 1951). Furthermore, subliminal processing could be observed using EEG (Kiss and Eimer, 2008; Kongthong et al., 2013).

Unconscious processing of stimuli could be facilitated through the use of highly arousing stimuli, and more specifically, fearful stimuli, given the evolutionary advantage of such mechanism and the existence of a fast pathway for fear in the human amygdala (Méndez-Bértolo et al., 2016). Examples of such stimuli would be snakes or spiders, commonly used to achieve fearful responses (Mayer and Merckelbach, 1999; Courtney et al., 2010; Lipka et al., 2011).

However, most of the research done in the field of subliminal stimulation has been done under laboratory settings, with few experiments having been carried out under more ecologically valid conditions (Boag, 2009).

5.1.3 Subliminal Stimulation in an Ecologically Valid Navigation Task

A study that addressed the use of unconscious processes for improving the human-computer interaction did so by using subliminal stimuli in a navigation task (Cetnarski et al., 2014). One of the main goals was to generalize the outcome of previous studies on subliminal perception, which have traditionally been conducted under controlled laboratory conditions, in a more ecologically valid environment. In order to achieve this, the eXperience Induction Machine (XIM) was used (Figure 5.1A). The XIM is an immersive environment that allows for psychophysiological experimentation under ecologically valid conditions in mixed and virtual reality (Bernardet et al., 2010; Betella et al., 2014b).

The results of that study confirmed that, indeed, subliminal stimuli can bias decision-making unconsciously, and effectively improve user perfor-

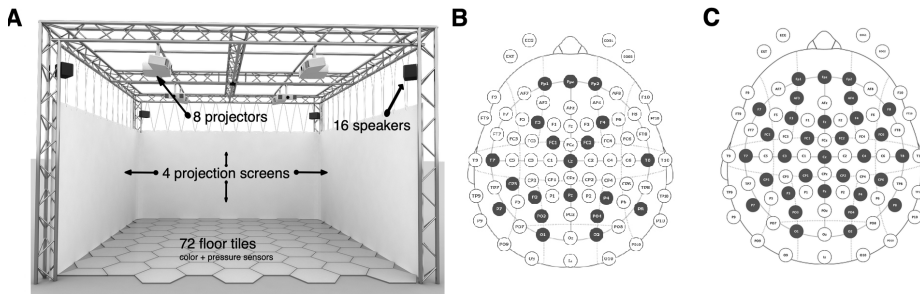


Figure 5.1: **(A)** Schematic illustration of the eXperience Induction Machine (XIM), an immersive space for mixed and virtual reality applications used to conduct experiments. It covers a space of 5.5 by 5.5 m equipped with multiple sensors (cameras, pressure-sensitive floor, etc.) and effectors (surrounding projectors, speakers, etc.). **(B)** EEG montage used in Experiment 1, with 20 electrodes, following the 10-20 international system. **(C)** EEG montage used in Experiment 2, with 32 electrodes, following the 10-20 international system.

mance on the navigation task (Cetnarski et al., 2014).

In order to obtain a mechanistic understanding of the underlying processes, we extended the previous study by simultaneously recording electrophysiological data during the navigational task. EEG data from the scalp have been used to observe the processing depth of the subliminal stimuli by pinpointing the brain structures that are activated. The second improvement was to try selecting the most effective stimulus for each user in order to optimize the effect of the subliminal aversive stimuli. Finally, to further improve the ecological validity of the study, the interaction was not restricted to the use of a keyboard but it was mapped to a full body interaction paradigm.

5.2 Materials and Methods

We carried out two different experiments, the second one being a modified version of the first in order to correct issues by modifying the navigation

and to further explore the neural activity by using a higher resolution EEG system and a condition for higher cognitive load.

In order to capture the different psychophysiological signals, we employed a novel framework that we developed, based on lab streaming layer (LSL ¹), a system for integrating and synchronizing signals from different sources. We needed such a framework, designed as a distributed architecture, in order to synchronize signals from multiple sensing devices (such as EDA sensors, EEG, and eye tracker), connected to different computers of the XIM, the immersive space where our study was carried out, while also recording event markers. Our architecture allows for a time-precise synchronized recording of all of these signals, integrated in a mixed and virtual reality space. Additionally, we deployed the MIDAS (Modular Integrated Distributed Analysis System) framework (Henelius and Tornaiainen, 2018), which allows us to compute online analysis of the acquired signals for near-real-time usage in the virtual experience for augmented interaction. In the present study, we did not use this component of the architecture except for in a pilot experiment.

5.2.1 Experiment 1

Participants

12 voluntary subjects (5 females and 7 males, mean age = 25.7 yrs, $SD = \pm 2.9$) recruited from the university campus participated in the study. EEG data of 3 subjects were discarded for analysis of the calibration phase due to technical issues. All participants reported normal or corrected-to-normal vision. All the subjects read and signed an informed consent form declaring that they clearly understood all the experimental procedures and the aim of the study. The study was approved by the local Ethical Committee (CIREP-UPF).

¹<https://github.com/sccn/labstreaminglayer>

Acquisition Sensors

EEG data were recorded with an Enobio20 wireless system (Neuroelectronics, Spain) at 500 Hz using 20 Ag/Ag-Cl electrodes placed accordingly to the international 10-20 system (Figure 5.1B), whereas subject's movements were recorded through a 3 axial accelerometer embedded in the EEG controller. For the electrodermal data, a BITalino device (Plux Wireless Biosignals S.A., Portugal) was used using two pre-gelled electrodes placed on the palm of the left hand. A Pupil eye tracker (Pupil Labs UG, Germany) was used to track the size of the pupil and eye-gaze. Signals from the different sensors were synchronized using the LSL protocol with LabRecorder (SCCN, University of San Diego, CA, USA). Due to synchronization issues, eye's data have been excluded from the analysis. Finally, the virtual maze environment was developed using the Unity 3D engine (Unity Technologies, San Francisco, CA, USA).

Experimental Procedure

The main experiment consisted of a navigation task in a 3D virtual maze. In this virtual maze, the user's point of view advances forward automatically through a straight, dimly-lit corridor until arrival at a vertical bifurcation. This Y-junction is established with a door on the ground. At this point, the stimulus sequence is displayed, showing the fixation point (500 ms), a subliminal prime (16 ms), a mask (500 ms) and a supraliminal target (4000 ms). The subject is then asked to take a decision: to either follow the current path (i.e., by leaning the head forward) or leave it by going through the door on the ground (i.e., by leaning the head backward). Finally, a feedback animation is provided, which can be positive (i.e., a gold ring acting as a reward) or negative (i.e., a 3D spider attacking the subject, which acts as a punishment). This sequence represents one full trial. After each trial, a new one begins until this main experimental block ends (150 trials).

An experimental block consists of two types of trials according to the supraliminal target displayed (see Figure 5.2): "fixed-choice" trials in which the supraliminal stimulus is a negative (a spider, different from the

one used as the subliminal prime) or a neutral image (an unrecognisable shape) informing the subject about the correct path to choose (i.e., avoid the current path in case of a spider) and “free-choice” trials in which the stimulus is represented by a question mark and the subject has to freely choose which path to follow. In fixed-choice trials, the accuracy of the response to the target (i.e., target compatibility of the response) determined the feedback type, whereas in free-choice trials the prime compatibility of the response determined the type of feedback.

Fixed-choice trials served the purpose of establishing the stimulus-response associations, whereas the “free-choice trials” serve as catch trials to study the behavioral impact of the subliminally perceived stimuli on a decision.

Before starting the main experiment a calibration phase took place aimed at selecting the most arousing stimuli for each participant. 20 stimuli (half neutral and half spiders images) obtained from the Geneva affective picture database (GAPED) (Dan-Glauser and Scherer, 2011) were randomly presented and each subject was asked to rate them through the Affective Slider (Betella and Verschure, 2016) while at the same time recording psychophysiological activity. Neutral images were selected among the ones with the lowest arousal score while the spider images were chosen among those with the highest arousal score to guarantee the largest difference between stimuli. All images were converted to grayscale, and the backgrounds of the spider images removed.

The protocol consisted of five main steps. Each participant was seated in the center of the XIM (about 2 meters away from the screen). It started with the reading and signature of the consent form, as well as the instructions for the experiment and the first anxiety questionnaire. After this, the participant was equipped with the sensors and the experimental session began with the calibration phase, followed by the navigation task. Once terminated, a visibility test was performed to ensure that the subliminally presented stimuli were indeed not consciously perceived. This visibility test placed the participants in the same scenario as in the main phase, but without movement. Participants were shown the fixation point, followed by the same subliminal images as in the main phase, and the masking.

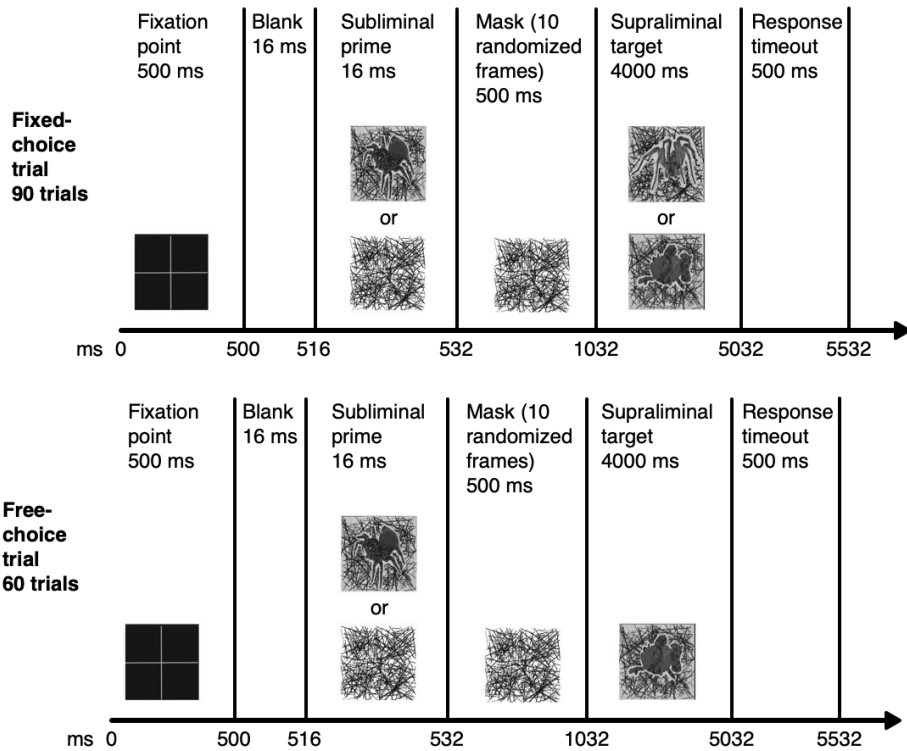


Figure 5.2: Stimulus sequence timeline showing the two types of possible trials in Experiment 1. In both fixed-choice trials and free-choice trials, a subliminal prime is shown. The difference is in the supraliminal target that is displayed, that either dictates the correct response (fixed-choice trial) or lets the subject decide (free-choice trial).

They were asked to indicate with a key press whether they saw a spider image or not, 40 times. Finally, the sensors were removed, the second anxiety questionnaire was administered, and the subject was debriefed.

Data Analysis

In order to compute the response accuracy rate during the main phase of the experiment, we computed the percentage of times that the prime-compatible response was made, i.e., the trial ended with a virtual gold ring, which means that when the negative prime was shown, the participant switched paths, and when the neutral prime was shown, the participant continued on the same path. This was computed by dividing the sum of all responses. The statistically significant deviation from chance level (50 %) represents the bias response. This was checked with a one-sample *t*-test. We also compared the accuracy depending on the stimulus type (negative or neutral) and compared it using a paired-samples *t*-test.

For the visibility test, the accuracy was computed in a similar fashion, by computing the percentage of times that the participants reported correctly to have seen a spider image. We checked the deviation from chance level (50 %) with a one-sample *t*-test. We computed the correlation between accuracy during the visibility test and during the main phase using a Pearson correlation test.

In order to analyze the EEG data, the EEGLAB (Delorme and Makeig, 2004) and MNE (Gramfort et al., 2013, 2014) software libraries were used. Ledalab (Benedek and Kaernbach, 2010) was used for EDA data analysis, following a continuous decomposition analysis (CDA). Particularly, we computed the sum of SCR-amplitudes of significant SCRs within the response window (reconvolved from corresponding phasic driver-peaks) and the average phasic driver within the response window. We separated these results depending on the participants movement choice (to avoid or not the virtual spider attack based on the subliminal cue) and compared them using a Mann-Whitney *U* test.

5.2.2 Experiment 2

Participants

16 voluntary subjects (9 females and 7 males, mean age = 24.5 yrs, $SD = \pm 3.2$) recruited from the university campus participated in the study. One subject was excluded due to a technical issue with the interaction device. Another subject was excluded because he reported being able to see the subliminal stimuli. All participants reported normal or corrected-to-normal vision. All the subjects read and signed an informed consent form declaring that they clearly understood all the experimental procedures and the aim of the study. The study was approved by the local Ethical Committee (CIREF-UPF).

Acquisition Sensors

EEG data were recorded with an Enobio32 wireless system (Neuroelectronics, Spain) at 500 Hz using 32 Ag/Ag-Cl electrodes placed accordingly to the international 10-20 system (Figure 5.1C). For the electrodermal data, an e-Health Sensor Shield (Libelium, Spain) was used. A Pupil eye tracker (Pupil Labs UG, Germany) was used to track the size of the pupil and eye-gaze. Signals from the different sensors were synchronized using the LSL protocol with LabRecorder (SCCN, University of San Diego, CA, USA). Finally, the virtual maze environment was developed using the Unity 3D engine (Unity Technologies, San Francisco, CA, USA).

Experimental Procedure

In Experiment 2, the virtual maze was modified to have a horizontal bifurcation (left/right decision) instead of a vertical one in order to avoid having a default movement direction (continue forward in Experiment 1). Due to this new virtual space configuration, a new stimuli exposure setup was needed. In Experiment 2, two stimuli were presented side by side simultaneously. The negative stimulus was presented randomly on either the left or the right side, with the neutral stimulus on the other side.

The interaction method was changed from head movements to remote controllers in order to avoid the movement biases found during Experiment 1. In Experiment 2, participants had identical wireless controllers (Pyrus PY-1) in both hands. They were instructed to press the trigger of the left controller to choose to move left and the trigger of the right one to move right.

Trial duration and distribution was modified to increase the number of trials, particularly of free-choice trials. In Experiment 2, there were 30 fixed-choice trials and 160 free-choice trials, each trial lasting from approximately 6 seconds (Figure 5.3). Fixed-choice trials were distributed non-randomly: 15 at the beginning of the experiment in order to establish the stimulus-response association and 1 every 10 free-choice trials.

A high cognitive load condition was introduced. While in the control group the score was shown at all moments, in the high cognitive load group participants had to keep mental track of the score and report it every 10 trials. Trials that ended with a gold ring awarded 3 points, while the ones that ended in a spider attack subtracted 5 points.

The calibration phase of Experiment 1 was removed. In Experiment 2, the negative stimulus (spider image) was the same for all subjects. This image was selected as the one that induced the higher arousal response in Experiment 1 according to the reported ratings.

Data Analysis

The procedure to compute and test the response accuracy during the main phase and the visibility test were the same as in Experiment 1, aggregated for both conditions (higher cognitive workload and lower cognitive workload). Additionally, we compared the response accuracy rate between both conditions by separating both groups and using an independent-samples *t*-test.

As in Experiment 1, EEGLAB and MNE were used to analyze EEG data. Ledalab was again used for EDA data analysis, computing and comparing the same two main measures.

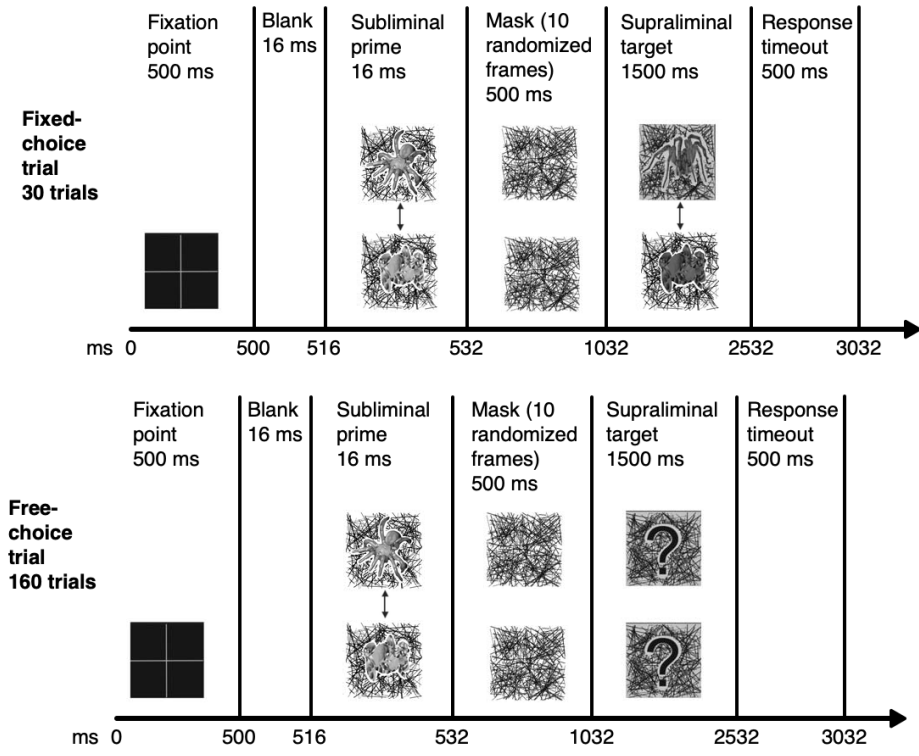


Figure 5.3: Stimulus sequence timeline showing the two types of possible trials in Experiment 2. In every trial, a negative subliminal stimulus and a neutral subliminal stimulus are shown, one of them on the right and the other on the left. As in Experiment 1, the difference between trial types is in the supraliminal target that is displayed, that either dictates the correct response (fixed-choice trial) or lets the subject decide (free-choice trial).

5.3 Results

5.3.1 Experiment 1

Behavioral Data

We started the statistical analysis with the visibility test data to confirm that the subliminal stimuli were indeed subliminal, and not consciously perceived, by most participants. 2 participants were removed from all analysis due to having an accuracy in this test above one standard deviation. Another participant was discarded as an outlier in response patterns. For the remaining participants, testing the visibility test scores (see Figure 5.4A) against chance level, there was no statistical significance, as expected, $t(9) = 0.751$, $p = 0.47$. Furthermore, there was no correlation between the scores in the visibility test and the accuracy in the main experiment, $r(9) = 0.440$, $p = 0.24$.

Accuracy scores were calculated as the proportion of prime congruent and incongruent responses of each subject in the free-choice trials, dividing the number of congruent choices by the sum of all responses. The mean accuracy score was 49.81 % (± 4.26 *SD*) (see Figure 5.4B) which was not significantly different from a random choice ($t(9) = -0.123$, $p = 0.91$). This result indicates that decisions were not biased by the prime presented.

Due to this mean accuracy, further analysis was performed on the data. First, by separating between trials with the negative prime and trials with the neutral prime, we observe a difference. The mean accuracy in trials with the negative prime is 37.91 %, while in trials with the neutral prime, it is 62.36 % (see Figure 5.4C). The difference in response accuracy between both types of trials was found to be significant, $t(9) = -3.6$, $p < 0.001$. This is a difference that was not expected.

Analyzing the frequency of forward/backward movements independently of stimulus kind (i.e., the decision to whether continue, by moving the head forward, or leave the path, by moving the head backward), a response pattern similar to the last one emerged. The subjects decided to move forward (continuing their path) 62.22 % of the time, and only

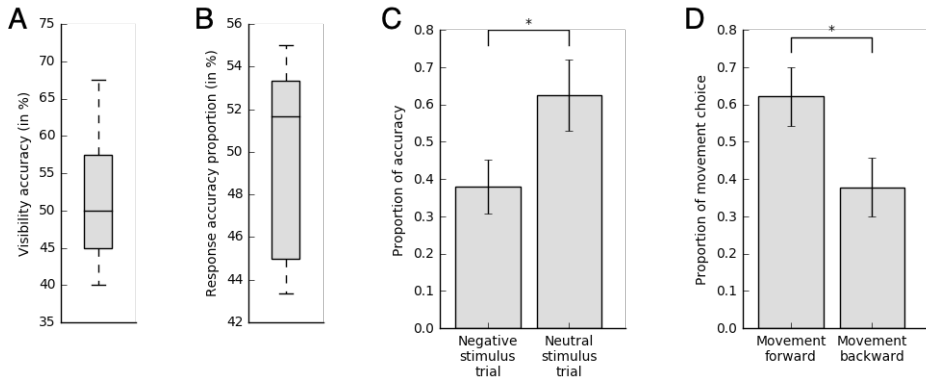


Figure 5.4: Behavioral results from Experiment 1. **(A)** Mean accuracy in the visibility test. This plot represents the mean accuracy percentage of the subjects during the visibility test. As expected, all values are around 50 %. **(B)** Average response accuracy in free-choice trials. A 50% value represents the chance level for a binary task. **(C)** Mean accuracy in free-choice trials, separated by stimulus type. This plot represents the mean accuracy for the subjects, taking into account the free-choice trials of the main experiment, separated between trials in which the negative prime was shown, and those in which the neutral prime was shown. The error bars represent the 95% confidence interval. **(D)** Percentage of movement that was performed in free-choice trials. This plot represents the percentage of movements that were performed during the main experiment, considering only the free-choice trials. A preference for moving forward can be observed. The error bars represent the 95% confidence interval.

37.78 % backward (leave their path; see Figure 5.4D). This difference was statistically significant, $t(9) = 3.593$, $p < 0.001$. Again, this is an effect that was not expected, as movements forward and backward should have been distributed 50/50 approximately. This is true both if we only consider chance level (no priming) and if we consider the priming effect, given that there was the same number of both types of stimuli.

It is important to remember that, when the neutral prime is presented, the correct response (the one that leads to the reward) is to move forward, and when the negative prime is presented, the correct response is to move backward. Thus, this preference for moving forward can explain the higher accuracy with the neutral prime and lower accuracy with the negative prime.

The analysis of the response times did not reveal a significant priming effect. Difference regarding the subliminal stimulus that was presented did not reach statistical significance, $t(9) = -1.816$, $p = 0.11$. There was no correlation between response time and accuracy, $r(9) = -0.344$, $p = 0.365$.

Physiological Data

The analysis of the EEG data revealed a significant effect of the images that were shown to the subjects. By taking into account the amplitude of the raw signal, it is possible to differentiate the brain responses between neutral and spider images, both when presented supraliminally (in the calibration phase), $t(180) = 5202.0$, $p < 0.01$, and when presented subliminally (in the main phase), $t(180) = 6234.0$, $p < 0.01$.

Focusing on the brain areas that were activated during the presentation of the stimuli, we can indeed observe that the frontal lobe presented differences in activation depending on the type of subliminal stimuli, $t(180) = 87.0$, $p < 0.05$. However, this difference was not detected in the occipital lobe (more specifically, in the visual cortex), $t(180) = 48.0$, $p = 0.10$.

Regarding the arousal, it was also possible to differentiate between neutral and spider images. When presented supraliminally (in the calibra-

tion phase), the best indicator was the alpha power (or beta power, which produced very similar results), especially in the frontal area, $t(90) = 1244.0$, $p < 0.01$. When presented subliminally (in the main experiment), the best indicator was the beta/alpha ratio, in the whole brain, $t(177) = 16099.5$, $p < 0.05$.

These arousal results were validated with those computed with the EDA signal. For computing the arousal with EDA data, the signal was filtered with a median filter to get the phasic data. The result is a significant correlation between both, $r(80) = 0.322$, $p < 0.01$. However, there was no significant correlation between both and the slider, $r(180) = 0.075$, $p = 0.315$, which might be due to insufficient quality of the EDA signal that was acquired.

During the main phase, we found differences in electrodermal response on the moment the subliminal stimulus was presented depending on the eventual movement choice of the participants, which happens later. Considering the sum of SCR-amplitudes of significant SCRs within the response window (reconvolved from corresponding phasic driver-peaks), there was a difference in electrodermal response according to the trial outcome ($U(1200) = 122400.5$, $p < 0.05$). There was also a difference in this when considering the average phasic driver within the response window ($U(1200) = 122480.0$, $p < 0.05$).

5.3.2 Experiment 2

Behavioral Data

Like in Experiment 1, we started the statistical analysis with the visibility test data to confirm that the subliminal stimuli were indeed subliminal, and not consciously perceived, by most participants. 2 participants were removed from all analysis due to having an accuracy in this test above one standard deviation. For the remaining participants, testing the visibility test scores (see Figure 5.5A) against chance level, there was no statistical significance, as expected, $t(12) = 0.491$, $p = 0.63$. Furthermore, there was no correlation between the scores in the visibility test and the accuracy in the main experiment, $r(12) = 0.205$, $p = 0.523$.

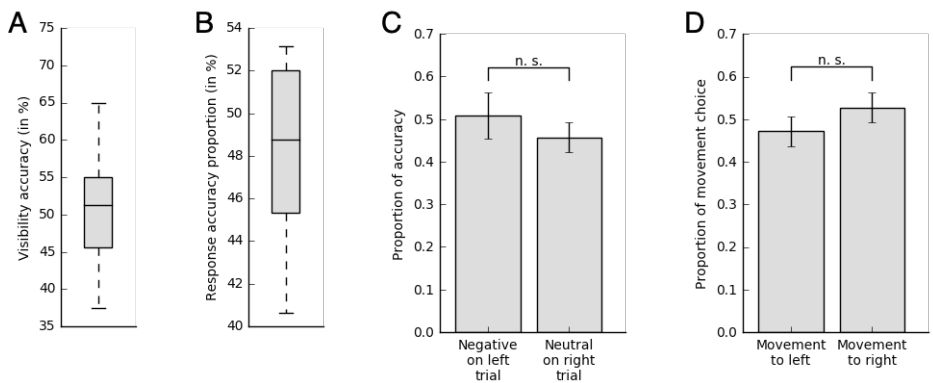


Figure 5.5: Behavioral results from Experiment 2. **(A)** Mean accuracy in the visibility test. This plot represents the mean accuracy percentage of the subjects during the visibility test. As expected, all values are around 50%. **(B)** Average response accuracy in free-choice trials. A 50% value represents the chance level for a binary task. **(C)** Mean accuracy in free-choice trials, separated by stimulus side. This plot represents the mean accuracy for the subjects, taking into account the free-choice trials of the main experiment, separated between trials in which the negative prime was shown on the left or on the right. There is no significant difference between both sides. The error bars represent the 95% confidence interval. **(D)** Percentage of movement that was performed in free-choice trials. This plot represents the percentage of movements that were performed during the main experiment, considering only the free-choice trials. No significant preference for moving left or right can be observed. There is no significant bias towards any direction. The error bars represent the 95% confidence interval.

Accuracy scores were calculated as the proportion of prime congruent and incongruent responses of each subject in the free-choice trials, dividing the number of congruent choices by the sum of all responses. The mean accuracy score for all subjects aggregated was 48.38 % (± 4.06 *SD*) (see Figure 5.5B) which was not significantly different from a random choice ($t(12) = -1.318, p = 0.214$). This result indicates that decisions were not biased by the prime presented in a consistent way.

Although the response accuracy was several percentage points below 50 % for half of the subjects, it was above this threshold for the other half, suggesting that the effect of the subliminal priming could be different between subjects.

The accuracy for the cognitive workload group was 48.75 % (± 4.75 *SD*) and 47.88 % (± 2.76 *SD*) for the non-cognitive workload group. No significant differences in accuracy scores were found between both groups ($t(12) = 0.338, p = 0.743$).

Regarding the side on which the negative stimulus was presented, the accuracy was 51.20 % when it was presented on the left side, and 45.83 % when it was presented on the right side (see Figure 5.5C). Although no statistical difference was found between both cases ($t(12) = 1.522, p = 0.156$), the response accuracy when the negative stimulus was shown on the right was significantly different from chance level ($t(12) = -2.67, p = 0.022$). This was not the case when the negative stimulus was presented on the left ($t(12) = 0.308, p = 0.764$). No significant navigation bias for either left or right movements was found (see Figure 5.5D) ($t(12) = -1.733, p = 0.111$). 100 % of the subjects were right-handed.

There was no significant statistical difference in response time depending on whether the response was correct (the spider was avoided) or not ($t(1920) = 450494.0, p = 0.210$), with a mean response time of 0.560 seconds (± 0.181 *SD*) when the trial was correct versus 0.567 seconds (± 0.210 *SD*) when it was incorrect. Additionally, there was no correlation between average response time and accuracy proportion ($r(12) = -0.957, p = 0.957$).

Physiological Data

As in Experiment 1, there was a difference in electrodermal response on stimulus presentation depending on the choice that participants made later. Considering the sum of SCR-amplitudes of significant SCRs within the response window (reconvolved from corresponding phasic driver-peaks), there was a difference in electrodermal response according to the trial outcome ($U(2850) = 921182.5, p < 0.001$). There was also a difference in this when considering the average phasic driver within the response window ($U(2850) = 909068.0, p < 0.001$). There was no significant difference in electrodermal response between the first and last 25 % of trials ($U(1200) = 173614.5, p = 0.142$).

There was a significant difference in electrodermal response between the cognitive workload group and the non-cognitive workload group. Considering the the average phasic driver within the response window of all trials, there was a difference both when taking all responses together ($U(1900) = 324205.5, p < 0.001$), and when considering separately responses that ended in a spider attack ($U(890) = 68413.0, p < 0.001$) and in a gold ring ($U(1010) = 88603.0, p < 0.001$).

A significant increase in spectral power was observed within the high delta band for the negative primes in the fronto-parietal-occipital electrodes when compared to the neutral primes. A decrease in high theta/low alpha was instead observed in the right occipital regions. Beta/alpha power ratio in the frontal electrodes appears to be a good predictor of arousal level induced by subliminal primes.

5.4 Discussion and Conclusion

In this study, we have investigated the behavioral and physiological effects of the presentation of subliminal emotionally charged stimuli presented on an ecologically valid navigation task in an immersive virtual-reality environment. To do this, we carried out two sequential experiments, based on a previous study (Cetnarski et al., 2014), in which participants had to navigate through a virtual maze while being exposed to

subliminal stimuli (a spider image and a neutral shape) on bifurcation points (forward or downward on Experiment 1 and left or right on Experiment 2). We expected to bias the navigation decisions of the participants by priming them to avoid certain paths using the subliminal stimuli as cues, as in the previous study. Although the subliminal priming did not consistently affect the navigation decisions of the participants in our experiments, we found significant and relevant results in the physiological responses.

Regarding the behavioral responses of the participants, we observed a significant bias to move forward in Experiment 1, most likely due to the navigation paradigm that was used, through head movements, as revealed from the data. We solved this in Experiment 2 by changing the design of the virtual environment to a left/right navigation decision selected using remote controllers. Still, we observed an inconsistent response accuracy across participants, with the mean not significantly different from chance level. These results, along with the distribution of response accuracy, lead us to speculate that the short subliminal stimulus presentation (16 ms) did not allow for the emotional charge of the images to be consistently processed.

Even if this is the case, we confirmed that the subliminally presented stimuli were indeed having an effect, as demonstrated by both EEG and electrodermal activity results. In both experiments, we observed a significantly different response in EDA locked to the subliminal stimulus presentation regarding the navigational choice that would be performed by participants moments later.

In conclusion, this project highlights the potential for using electrophysiological measurements such as EEG in tasks that present naturalistic experimental environments thanks to virtual reality. We have observed how the subconscious reactions of the participants can be tracked using EEG, even with extremely short exposure times to the stimuli. An extended version of the framework for psychophysiological sensing that we used could be employed in future experiments in order to perform online analysis of the acquired EEG signal for implicit real-time interaction with the virtual reality environment.

Chapter 6

KEYSTROKE DYNAMICS CORRELATE WITH AFFECTIVE CONTENT

This chapter is based on:

López-Carral, H., Santos-Pata, D., Zucca, R., and Verschure, P. F. (2019). How you type is what you type: Keystroke dynamics correlate with affective content. *2019 8th International Conference on Affective Computing and Intelligent Interaction, ACII 2019*, pages 359–363

Estimating the affective state of a user during a computer task traditionally relies on either subjective reports or analysis of physiological signals, facial expressions, and other measures. These methods have known limitations, can be intrusive and may require specialized equipment. An alternative would be employing a ubiquitous device of everyday use such as a standard keyboard. Here we investigate if we can infer the emotional state of a user by analyzing their typing patterns. To test this hypothesis, we asked 400 participants to caption a set of emotionally charged images taken from a standard database with known ratings of arousal and valence. We computed different keystroke pattern dynamics, including

keystroke duration (dwell time) and latency (flight time). By computing the mean value of all of these features for each image, we found a statistically significant negative correlation between dwell times and valence, and between flight times and arousal. These results highlight the potential of using keystroke dynamics to estimate the affective state of a user in a non-obtrusive way and without the need for specialized devices.

6.1 Introduction

Estimating the affective state of users while interacting with computers attracted much interest in recent years due to its potential for enhancing Human-Computer Interaction (HCI). This has been the focus of the affective computing field for the last few decades, with the expectation of having an impact in many fields that depend —increasingly so— in HCI, such as education, robotics, or human health, among others. Furthermore, progress in affective computing could also help to advance our knowledge of emotions and human cognition (Picard, 1995).

Typical experimental approaches for inferring affective states include subjective reports, in which users are repeatedly asked to describe or rate how they are feeling, or estimates from measures such as physiological responses, facial expressions or body gestures. Subjective reports are a traditional tool in a large portion of affective sciences. Multi-item scales (Watson et al., 1988) or pictorial tools (Bradley and Lang, 1994) can be easily administered both in paper or digital formats. Meanwhile, estimations from bodily responses or expressions require devices that have to be either worn by the users or placed close to them. Examples of this include inferring emotional states by detecting changes in electrodermal activity (EDA) (Critchley, 2002; Lang et al., 1993) and heart rate variability (HRV) (Agrafioti et al., 2012; Selvaraj et al., 2013). Facial expression analysis typically requires a classification between a set of discrete basic emotions (Fasel and Luetten, 2003).

These methods, although valid in assessing affective states, present several issues. Subjective reports interrupt the regular user’s workflow

and present known limitations in terms of validity and reliability (Stone et al., 1999; Keefer, 2015). In the case of bodily responses, different devices are required to be working alongside the computer, in many cases in direct physical contact with the user (e.g., electrodes placed on the skin), which can be intrusive to the user and expensive due to the economic cost of these devices.

A way to overcome these problems would be using components that are already available when interacting with a computer. Devices that do not require any unusual or particular action from the user, while still able to provide relevant correlates of internal states. Several studies have previously attempted to infer specific user's traits by analyzing the way a person types on a keyboard. The initial observation of unique typing rhythms across individuals (Joyce and Gupta, 1990) promoted in the last three decades great interest in the study of keystroke dynamics, particularly in the field of user authentication (Teh et al., 2013). Several studies have succeeded in authenticating users with high accuracy based on a variety of classification algorithms (Monrose and Rubin, 1997; Bergadano et al., 2003), suggesting that the relevant component in typing is not only the content typed but also how it is typed (Monrose and Rubin, 2000).

More recently, some studies have also provided evidence of the possibility of using keystroke dynamics to estimate emotional states (Kolakowska, 2013). In a recent field study, typing rhythms were coupled with periodic self-reports to classify between a series of discrete emotional states (Epp et al., 2011) with high accuracy. In a different study, keystroke patterns were successfully used to estimate the level of individual stress (induced by a mental arithmetic test), as self-reported in a pre- and post-questionnaire, as well as heart rate variability (Gunawardhane et al., 2013).

To the best of the authors' knowledge, no previous study has investigated the possibility of adopting keystroke dynamics to discriminate the affective features of presented stimuli that participants are describing. This would imply that the emotional content of these stimuli is actively affecting the participant's way of typing in subtle ways that could be detected through the analysis of their typing patterns. In this study, we show

that this is indeed a possibility. To do this, we asked 400 participants to caption a set of images rated in terms of continuous values of arousal and valence, instead of discrete emotions. The arousal and valence dimensions of affect follow the circumplex model of affect (Russell, 1980), which is generally used in experiments related to affect induction and detection (Gunes et al., 2011).

6.2 Methods

We performed an online study where participants ($N = 400$) were asked to observe a series of 46 images and to type a description for each one of them.

6.2.1 Participants

Participants were recruited using the Amazon Mechanical Turk (MTurk) service. To guarantee high-quality responses, we restricted participation to volunteers with at least 50 tasks previously completed and an approval rate of over 90 % and with a proficient level of English. No personal information from the participants (such as names or IP addresses) was recorded. As expected from this pool of participants, there was considerable variability in demographic backgrounds, with a mean age of 37.51 ($SD = 12.17$). 60.4 % of participants were male. 47.11 % reported spending over 5 hours per day typing on a keyboard, 38.35 % between 3 and 5 hours, and 14.54 % fewer hours.

6.2.2 Experimental Protocol

Before starting the task, a page informed the participants about the experimental protocol and their right to recess at any moment. Subsequently, they were required to answer a set of questions including demographics (age, gender, education level, and primary language), keyboard experience (number of hours per day spent typing) and keyboard layout used. After this, the main task started.

The main task consisted of visualizing a sequence of 46 images and providing a description of each one. In each trial, an image was presented for 2 seconds, and it was followed by a text field in which the participants were instructed to type a free description of the image seen using a minimum of 4 words.

The images were selected from the Open Affective Standardized Image Set (OASIS) (Kurdi et al., 2017), a set of 900 open-access images available for online use with normative ratings of arousal and valence following Russell’s circumplex model of affect (Russell, 1980). From the whole set, 46 images were extracted to cover the entire range of arousal and valence ratings. To select the images, we first binned the dataset into a 6 by 6 matrix along the two dimensions of valence and arousal (OASIS ratings are expressed using a 7-point Likert scale), and we selected two images at random from the 23 bins which contained at least 2 images. Thus, 46 images were finally selected (see Figure 6.1 which were randomly presented to every participant).

6.2.3 Data Collection and Processing

During the task session, all keyboard events performed on the provided text field have been logged. Specifically, this includes both key-down (or key-press) and key-up (or key-release) events (see Figure 6.2). For each event, three pieces of information were stored: the type of event (key up or down), the key that was pressed, and the timestamp (in milliseconds). We also recorded the time in which each image and the text field were presented to the participants, as well as when they clicked to proceed to the next image. All these data were collected in a file for each participant in JSON format. The mean duration of the experiment was 14.52 minutes, and the mean length of the image descriptions was 5.93 words.

After collection, data were inspected to ensure that there were no instances of image descriptions that were copied and pasted, missing data due to connectivity issues or other problems that might affect the integrity of the acquired data. To do this, we reconstructed each final image descriptions provided by the participants from the individual keystroke




Figure 6.1: The selection process of the 46 images included in the experimental dataset. Two images were extracted at random from each of the 23 non-empty bins in which the 900 OASIS image set was previously divided (Kurdi et al., 2017). This selection covers the entire range of arousal and valence. Dark circles represent the selected images, while semitransparent ones represent the rest of the images in the original set.

events.

For each image description provided by the participants, we derived a series of features. Two standard features in keystroke dynamics are the duration and the latency (Monrose and Rubin, 2000; Bergadano et al., 2003; Epp et al., 2011). Keystroke duration (also known as *dwel time*) represents the time that a single key was pressed in an instance (time since key-down until key-up). Keystroke latency (or *flight time*) represents the elapsed time between two sequential key presses (time since key-up until next key-down). Additionally, we computed the number of error corrections (presses of the backspace key), the total time to write each descrip-




Figure 6.2: Visualization of keystrokes during an example image description. An example of the graphical representation of one instance of flight time and dwell time is provided in the annotations.

tion, the time since the presentation of the text field until the participants started typing, and the time since the participants finished typing until they pressed the button to continue. See Table 6.1 for a summary.

For each of the derived features, we computed the Spearman’s rank correlation coefficient between each feature and the reported arousal and valence ratings.

6.3 Results

In order to assess how affective properties of perceived content modulate typing behavior, we tested the interplay between the previously described features (see Table 6.1) and the affective ratings of the stimuli as provided

Table 6.1: Features extracted from participant’s behavior while typing

Feature	Description
Dwell time	Keystroke duration, time a key was pressed
Flight time	Keystroke latency, time between two key presses
Backspace count	Number of error corrections
Time total	Time since text field was available until submission
Time start	Time to start typing since text field appears
Time end	Time since finishing typing until submission

in the image set.

We first asked how valence affects the typical typing behavior of individuals. To do so, we extracted each participant average flight times and grouped participant’s scores per valence binned category. Statistical testing revealed a significant negative correlation between valence and typing flight times, with shorter flight times consistently associated to high valence ratings (Figure 6.3, top-left, Spearman-test $r = -1.0$, $p < 0.001$). A similar analysis was then performed for each participant average of dwell times and again grouped the individuals’ scores per valence binned category. Statistical testing revealed a significant interplay between valence and typing dwell times (Figure 6.3, top-middle, Spearman-test $r = -1.0$, $p < 0.001$).

Next, we examined whether the content arousal score would reveal similar effects in the typing behavior of individuals.

As for arousal analysis, we extracted each participant average flight times and grouped the individuals’ scores per arousal binned category. Statistical testing revealed a significant interplay between arousal and typing flight times (Figure 6.3, bottom-left, Spearman-test $r = -1.0$, $p < 0.001$). Similarly, we analyzed each participant average of dwell times and the individuals’ scores per arousal binned category. Statistical testing revealed a significant interplay between arousal and typing dwell times (Figure 6.3, bottom-middle, Spearman-test $r = -1.0$, $p < 0.001$).




Figure 6.3: Correlation analysis of each of the extracted features with valence and arousal considering bins of these affective ratings. The three features of interest are shown. We can observe highly significant correlations between both flight times and dwell times with both arousal and valence. Additionally, there is an highly significant correlation between the time to start typing since the text field was presented (Times start) and valence. There is no significant correlation between the start time and arousal ($r = -0.7$, $p = 0.188$). Times are expressed in milliseconds.

So far we have reported how affective features modulate the typing

behavior of individuals. However, in a natural environment, the valence rate of a stimulus does affect how quickly humans, and other animals, react to that same stimulus. To test whether, in a goal-oriented typing task, the onset of typing could be predicted by the affective content upon which participants were reporting, we grouped the individuals time to start typing scores accordingly to the stimuli affective rate. We observed a significant negative correlation between the onset of typing (initiation of behavior) and valence rate (Figure 6.3, top-right, Spearman-test $r = -1.0, p < 0.001$).

These metrics capture how much individuals modulate their typing profile based on the perceived and reported content. Next, we asked whether the content alone could generalize the participants typing behavior. To do so, we extracted the values of each feature individually for each image calculating the participant's population mean of the different features per image. We found a significant correlation between dwell times and valence ($r = -0.293, p < 0.001479$), as well as a highly significant correlation between flight times and arousal ($r = -0.377, p < 0.00109$) (see Figure 6.4). Therefore, suggesting an overall modulating effect between content and typing behavior.

6.4 Discussion and Conclusion

Being able to reliably estimate the affective state of a user while interacting with a computing device would significantly improve the interaction process, with machines that can be more reactive and adaptive. Such capability could be beneficial for diverse fields such as education, human-robot interaction, digital health, and others (Picard, 1995).

However, typical approaches for inferring these emotional traits rely on either subjective reports (e.g., (Bradley and Lang, 1994)) or on the usage of specialized equipment (e.g., (Critchley, 2002)), which can be intrusive and expensive. Therefore, a way to overcome these issues would be to use an automatic approach that would take advantage of a device that users would typically use, without interfering with their behavior.

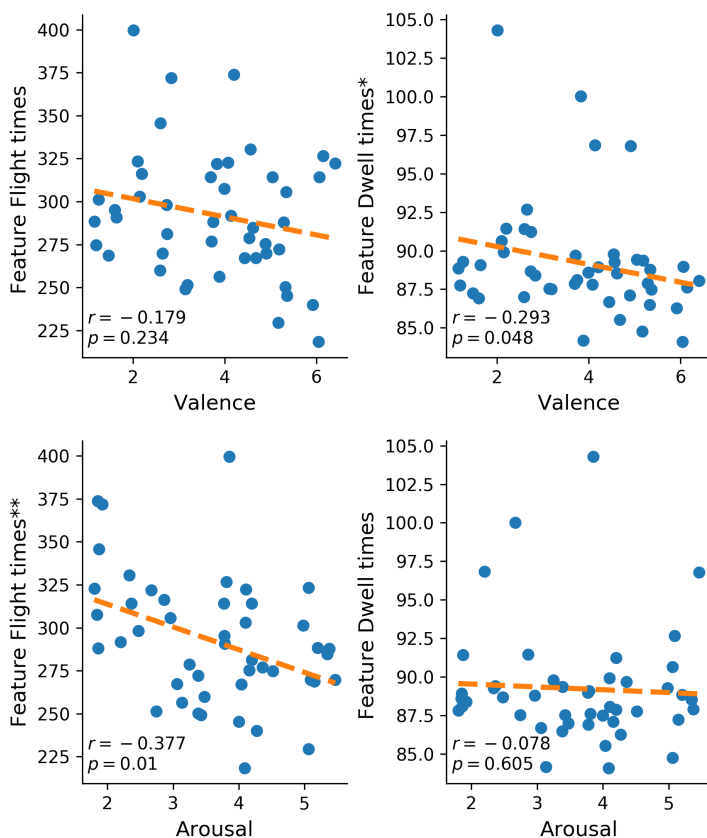


Figure 6.4: Correlation analysis of the mean value of the extracted features for each image with valence and arousal. Only the two features of interest are shown. We can observe a significant correlation between dwell times and valence, and between flight times and arousal. Times are expressed in milliseconds.

A possibility for this is to use keystroke dynamics computed from the typing patterns of users on regular keyboards. Such an approach has been explored mainly in the field of digital authentication (Monrose and Rubin, 2000), and only recently it has been extended to the field of affective

computing with promising results. However, current research within this field has relied still on subjective reports to validate the estimations or classifies emotional states within discrete states (Epp et al., 2011).

In this study, we analyzed the typing patterns of a large sample of participants that were asked to describe a set of images selected from the OASIS normative database for affective research (Kurdi et al., 2017).

We processed the recorded data in order to extract a series of keystroke features, including keystroke latency and duration, and timings, for each participant and each image. Analyzing the keystroke dynamics of the participants, we found highly significant negative correlations of both flight times and dwell times with both arousal and valence, as well as between time to start and valence. We then checked for generalization on the content itself, finding significant negative correlations between dwell times and valence, and between flight times and arousal.

These results show that keystroke dynamics do indeed correlate with both arousal and valence. Therefore, it could be possible to infer affective states from keyboard activity.

Furthermore, we achieved these results by merely exposing participants to emotionally charged images (for 2 seconds each) and asking them to describe them, without using any subjective report, thanks to the fact that the images were already rated. Each participant rated only 46 images, which shows that not a lot of typing information is required from an individual user.

Although we have found significant results using a limited amount of keystroke features, it is possible that we could have obtained relevant results by using more sophisticated features such as digraphs (combinations of two letters) or trigraphs (combinations of three letters) (Dowland and Furnell, 2005; Epp et al., 2011), or by using a simultaneous combination of multiple features.

Our results highlight the potential for a more in-depth analysis. This could include sentiment analysis on the descriptions written by the participants, in order to test the correlation between those results and both the affective ratings provided in the image set and the keystroke features we computed. Furthermore, machine learning techniques could be em-

ployed to train a model capable of determining the affective state during the typing of a sentence by using the described keystroke features.

In conclusion, this study reveals the correlation between keystroke dynamics and affective content by using descriptions of images from a rated set. This showcases the potential of using keyboard activity in order to infer affective states, either in addition to other techniques (such as physiological signals) or as a replacement when they are not possible, with the benefit of being an unobtrusive and inexpensive method.

Chapter 7

SUBJECTIVE RATINGS OF EMOTIVE STIMULI PREDICT THE IMPACT OF THE COVID-19 QUARANTINE ON AFFECTIVE STATES

This chapter is based on:

López-Carral, H., Grechuta, K., and Verschure, P. F. (2020). Subjective ratings of emotive stimuli predict the impact of the COVID-19 quarantine on affective states. *PLoS ONE*, 15(8 August):1–15

The COVID-19 crisis resulted in a large proportion of the world's population having to employ social distancing measures and self-quarantine. Given that limiting social interaction impacts mental health, we assessed the effects of quarantine on emotive perception as a proxy of affective states. To this end, we conducted an online experiment whereby 112 participants provided affective ratings for a set of normative images and reported on their well-being during COVID-19 self-isolation. We found

that current valence ratings were significantly lower than the original ones from 2015. This negative shift correlated with key aspects of the personal situation during the confinement, including working and living status, and subjective well-being. These findings indicate that quarantine impacts mood negatively, resulting in a negatively biased perception of emotive stimuli. Moreover, our online assessment method shows its validity for large-scale population studies on the impact of COVID-19 related mitigation methods and well-being.

7.1 Introduction

In December 2019, Chinese health authorities reported a cluster of pneumonia cases in the city of Wuhan, in the Hubei province, caused by the novel coronavirus SARS-CoV-2 (COVID-19) (World Health Organization, 2020). By mid-March 2020, a total of 200,000 confirmed cases (The Center for Systems Science and Engineering, Johns Hopkins, 2020) had been reported worldwide, showing an exponential increase with the current number of identified cases exceeding 14 million, whereby Spain, Italy, and the United Kingdom are the most-affected European nations.

To prevent the spread of COVID-19, public health authorities have employed mitigation strategies and, in particular, quarantine (Centers for Disease Control and Prevention, 2017) and isolation, which are currently practiced across the globe. Mandatory mass quarantine restrictions, which include social distancing, stay-at-home rules, and limiting work-related travel outside the home (Rothstein et al., 2003) might impact both physical and mental health of the affected individuals (Nobles et al., 2020). Indeed, prolonged widespread lock-down and limiting social contact has resulted in post-traumatic stress disorder, depression, anxiety, mood dysregulations, and anxiety-induced insomnia during previous periods of quarantine (Miles, 2015; Brooks et al., 2020; Hossain et al., 2020). These, in turn, led to cognitive distortions and maladaptive behaviors, including suicide (Rubin and Wessely, 2020; Barbisch et al., 2015). A growing body of evidence from COVID-19 demonstrates that the current mass quarantine

has been producing similar adverse psychological effects, which might have long-lasting consequences on both individual subjects and society (Nobles et al., 2020; Holmes et al., 2020; Rajkumar, 2020; Torales et al., 2020). Moreover, it is unclear for how long and how frequent confinement measures will be put in place in the medium and long-term. Hence, understanding the specific impact of COVID-19 on mental health and the development of monitoring and diagnostic tools to identify individuals at risk are of critical importance.

Disturbances in mental health, including disorders of mood, are commonly assessed using explicit questionnaires and interview measures (Clark and Watson, 1991). Both clinician-rated and self-reported instruments have been used for decades (Smarr and Keefer, 2011). Some studies, however, have outlined noteworthy limitations of standard assessments of depression, such as conceptual and psychometric flaws (Bagby et al., 2004; Zimmerman et al., 2005; Gibbons et al., 1993; Bech et al., 1984; Maier et al., 1985). For instance, the Hamilton Depression Rating Scale (HDRS, (Hamilton, 1960)), which has been considered a gold standard in clinical practice as well as clinical trials, was widely criticized for its subjectivity as well as the multidimensional structure, which varies across studies hence preventing replication across samples as well as poor factorial and content validity (Bagby et al., 2004; Zimmerman et al., 2005; Gibbons et al., 1993; Bech et al., 1984; Maier et al., 1985; Fried and Nesse, 2015). Moreover, it is well-established that self-reports in psychological research can suffer from response bias such as socially desirable responding or a tendency to provide positive self-descriptions (Paulhus, 2002; Braun et al., 2001; Paulhus, 2017). To counteract possible response bias and suggestion effects, in the current study, we employed affective ratings of calibrated emotional stimuli as an implicit measure of mental state building on earlier validation studies of online emotional rating methods of calibrated emotional stimuli (Betella and Verschure, 2016).

Mood-state-dependent changes in emotional reactivity are reflected in emotion experience evaluations (Rottenberg et al., 2005). Indeed, there is converging evidence that ratings of affective stimuli might serve as a robust, indirect measure of mood. For example, empirical studies show

reduced subjective and expressive emotional responses to neutral and positive stimuli in depression, including in major depressive disorder (MDD) (Sloan et al., 1997, 2002; Dunn et al., 2004; Berenbaum and Oltmanns, 1992). Specifically, the results show significant negative shifts in emotional ratings of valence compared to the healthy controls such that patients judge the stimuli as substantially less pleasant. Alternatively, Borderline Personality Disorder (BPD) patients show hypersensitivity to emotional stimuli as compared to healthy controls (Bortolla et al., 2019). These findings support the notion that response to emotive stimuli is be altered in disorders of mood.

Given the mental health risk of medium to long-term isolation (Brooks et al., 2020; Hossain et al., 2020; Haney, 2003), it is relevant to develop methods that can effectively and unobtrusively assess and monitor the impact of the restriction of movement and social distancing on well-being and mental health. Hence, the goal of this study is to evaluate the effects of quarantine-induced changes in mood, as measured implicitly through the subjective ratings of emotional stimuli. We predicted that individuals in quarantine due to COVID-19 might present changes in their affective ratings that reflect their subjective experience of isolation. To test this hypothesis, we conducted an online experiment in which volunteers were asked to rate the affective content of a subset of standardized visual stimuli and report their current personal situation and experience related to the pandemic. We compared the affective ratings of valence (i.e., indicative of disturbances in mood) between groups of subjects in the pre-quarantine “normal” condition and under quarantine.

7.2 Materials and methods

7.2.1 Participants

After providing their consent, one hundred twelve subjects participated in the study (64.29 % females) with a mean age of 32.38 ($SD = 9.04$). The sample size of $N= 110$ was determined a priori using G*Power software version 3.1 (Kiel, Germany) based on $\alpha = 0.05$, power of 80%

and medium effect size (0.5). Volunteers accessed the online experiment using a URL (uniform resource locator) that was shared through social media and instant messaging platforms by the experimenters. 51.79 % of the subjects held postgraduate degrees or higher. Subjects originated from 19 different countries (30.36 % Spanish and 21.43 % Italian), and they lived in 17 countries (53.57 % in Spain and 16.07 % in Italy). This sampling approach was chosen to cover a range of countries that were similarly impacted by self-isolation measures. In particular, for the analyses, we included only those participants who were actively undergoing quarantine. Thus, all participants were uniform in their cultural traits (Gupta et al., 2002) and quarantine measures, including social isolation and distancing, the banning of social events and gatherings, the closure of schools, offices, and other facilities, and travel restrictions (Conti, 2020; Shah et al., 2020).

The reported data were collected between the 9th and the 20th of April 2020. The personal data of the subjects were anonymized and kept confidential. All participants were blind to the purpose of the study. Specifically, until the end of the session, subjects did not know the study's objective, which could bias their responses. However, they were informed about it at the end of the trial.

7.2.2 Materials

Affective Slider

We employed the Affective Slider tool (Betella and Verschure, 2016) for digital assessments of the arousal and pleasure dimensions of the emotive stimuli. Its design principles follow the circumplex model of emotion proposed by James Russell (Russell, 1979, 1980). In this bipolar model, arousal corresponds to the intensity of an affective response (i.e., evoked level of excitement), while valence represents the positivity or negativity of the response (i.e., happiness). Consequently, the Affective Slider consists of a pair of slider controls flanked by emoticons that correspond to the ratings of arousal and valence, respectively. Both sliders are oriented horizontally and located above each other (Fig 7.1). In this study, Af-

fective Slider served to allow for continuous subjective assessment of the presented images, thus counteracting methodological limitations of classical scales such as the Self-Assessment Manikin (SAM) (Bradley and Lang, 1994) especially when applied in online assessments (Betella and Verschure, 2016). During the experiment, the position of the two sliders on the screen (e.g., arousal on top of valence or vice versa) randomly changed at every trial to prevent the order-effects and automaticity in the responses.




Figure 7.1: Example of digital assessments of the arousal and pleasure using the Affective Slider (Betella and Verschure, 2016). On the left, there is an example image from the OASIS data set (Kurdi et al., 2017). On the right, there are the ratings. The top slider corresponds to arousal and the bottom one to valence. This visual order was randomized over trials.

Experimental stimuli: Open Affective Standardized Image Set (OASIS)

OASIS is a validated open-access data set, which consists of nine hundred images acquired online (Kurdi et al., 2017). Each stimulus includes normative ratings of both arousal and valence reported on a scale between 1 and 7 by 822 participants. The stimuli depict a variety of themes within

four categories that include people, animals, scenes, and objects. In contrast with the well-known International Affective Picture Set (IAPS) (Lang et al., 2008), OASIS allows for online use of the data set and provides more recent ratings. For the purpose of this study, we chose a subset of 30 images from the categories *people* and *scenes*, corresponding to 61.78 % of the entire set. The choice was determined by the content of the stimuli, which was related to social and outdoor activities. The subset was selected randomly from the whole set of images to achieve a representative sample (see Fig 7.2). The same set of 30 images was presented to all participants in a randomized order.




Figure 7.2: Distribution of the valence and arousal rating for the 30 images selected for this study (solid circles) and the OASIS data set of 900 images (semitransparent circles).

COVID-19 questionnaire

To evaluate the current personal and social situation of each participant and their subjective experience during the COVID-19 global health crisis, we created a custom questionnaire. The scale was composed of 14 items, including an optional field to provide personal comments related to the

quarantine period. The answers to the remaining questions were to be delivered using either a multiple-choice scale or standard sliders derived from the Affective Slider. In the case of the latter, subjects rated their level of agreement on a scale ranging from “not at all” to “very much”. The questionnaire was administered at the end of the experiment. For the analysis, we included only the data of those subjects who completed the questionnaire.

7.2.3 Procedure

The online experiment consisted of four main sections: (a) instructions, the consent form, disclaimer, as well as the collection of demographic data (gender, age, education level, country of origin, and country of residence), (b) experimental task, (c) COVID-19 questionnaire, and (d) explanation of the rationale of the study.

During the experimental task, each participant was presented with a sequence of thirty affective stimuli from the OASIS image set (Kurdi et al., 2017). Participants provided their ratings using the Affective Slider located on the right side of the image (Figure 7.1). Each stimulus remained visible until the submission of both ratings, which had no time limit, as in the experimental tasks of both the tool (Betella and Verschure, 2016) and the data set (Kurdi et al., 2017). Only when both ratings were provided, subjects could advance to the next image by clicking a separate button. After that, the next stimulus was immediately displayed together with the corresponding Affective Slider.

Once participants completed the experimental task, they were required to complete the COVID-19 questionnaire. Finally, after having submitted the questionnaire, participants were presented with a final page that included the experimental rationale and the researchers’ contact information.

7.2.4 Data analysis

Tests of normality were performed on the data, and subsequently, T-tests were used to identify differences between the affective ratings. All comparative analyses used two-tailed tests and a standard level of significance ($p < .05$). For each comparison, the effect sizes were computed using Cohen's d (Cohen, 1988). A Pearson product-moment correlation coefficient was computed for the subsequent linear correlation analyses. Fourteen participants who reported not being in quarantine were excluded from the analysis.

Finally, we applied machine learning techniques to evaluate the plausibility of predicting participants' personal situation and reported subjective state during the quarantine based on their valence ratings provided during the experiment. To achieve this, we trained a C-Support Vector Classification (SVC) model. Parameter tuning was performed using a grid search algorithm. The model was cross-validated to evaluate its performance based on the F-score. The classification was performed using the Scikit-learn machine learning library (Pedregosa et al., 2011).

7.3 Results

First, we assessed the linear relationship between the affective ratings of arousal and valence collected in the present experiment and those acquired in the original study (Kurdi et al., 2017). To this end, we computed the mean rating from all the subjects for each of the experimental stimuli and extracted the corresponding mean values from the OASIS data set. The analysis yielded high and significant positive correlation between the mean scores for both arousal ($r(30) = .77, p < .001$, see Fig 7.3A) and valence ($r(30) = .88, p < .001$, see Fig 7.3B).

Second, to test our hypothesis, we evaluated the existence of possible shifts in the affective ratings between the present study and the OASIS for the subsets of neutral and positive stimuli. In the neutral subset, we included all the images whose mean ratings for valence ranged between 3 and 5 ($N = 15$), while in the positive one, those whose mean valence

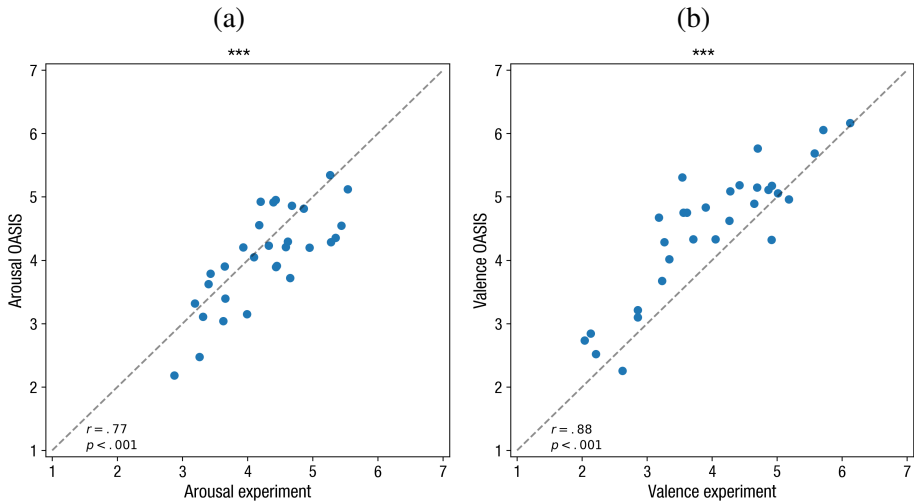


Figure 7.3: Linear correlations between the ratings obtained in our study and those from OASIS. A: Linear correlation between arousal ratings from OASIS (y-axis) and those acquired in the present study (x-axis). B: Linear correlation between valence ratings from OASIS (y-axis) and those acquired in the present study (x-axis). In both graphs, dashed lines represent the identity lines; * * * $p < .001$

ratings ranged between 5 and 7 ($N = 11$). For these analyses, we computed the mean rating of both arousal and valence from all subjects for each chosen subset. For the neutral stimuli, statistical analyses yielded that, while the mean ratings of arousal for the chosen images did not differ ($t(15) = .61, p = .546$), there was a statistically significant negative shift in the ratings of valence ($t(15) = -2.28, p = .030, d = .859$, see Fig 7.4).

Similarly, for the positive stimuli, we found no differences in the mean ratings of arousal ($t(11) = 1.313, p = .203$). In line with literature, however, we found a statistically significant negative shift in the ratings of valence ($t(11) = -2.148, p = .044, d = .974$).

Third, we conducted post hoc analyses to assess relationships between

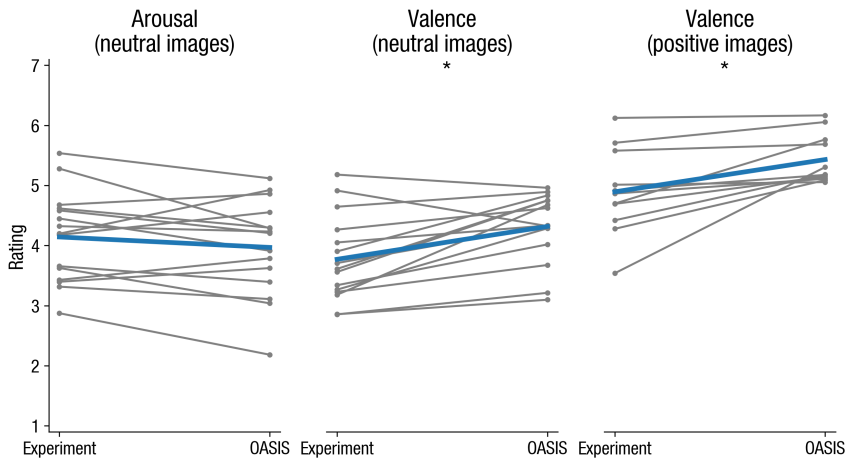


Figure 7.4: Shift in the affective ratings for neutral and positive images. The graphs present the comparison between the ratings of arousal (left) and valence for neutral images (middle) and valence for positive ones (right) obtained in our study with those from the OASIS. In all graphs, the blue lines correspond to the mean, while the individual lines show differences for individual images ($N = 15$); $*p < .05$

the affective ratings of valence and participants' situation during the quarantine period evaluated through the COVID-19 questionnaire. Specifically, we investigated if the mean ratings of valence are related to whether the subjects (a) enjoy working from home, (b) miss the “normal” pre-quarantine life, and (c) live alone. For these analyses, we computed the differences in mean ratings from the present study and the OASIS data set for each participant. The first correlation analysis yielded a significant positive linear relationship between the strength of the enjoyment of working from home and the mean difference in valence ratings ($r(98) = .24, p = .043$). In particular, we found that participants who reported enjoying working from home rated the images more positively than those who did not (Fig 7.5A). Second, our results revealed a significant negative correlation between the degree of missing the “normal” pre-quarantine life and the differences in valence ratings ($r(98) = -.22, p = .032$).

Hence, participants who missed more to return to the normal life rated images more negatively than those who missed it less (Fig 7.5B).

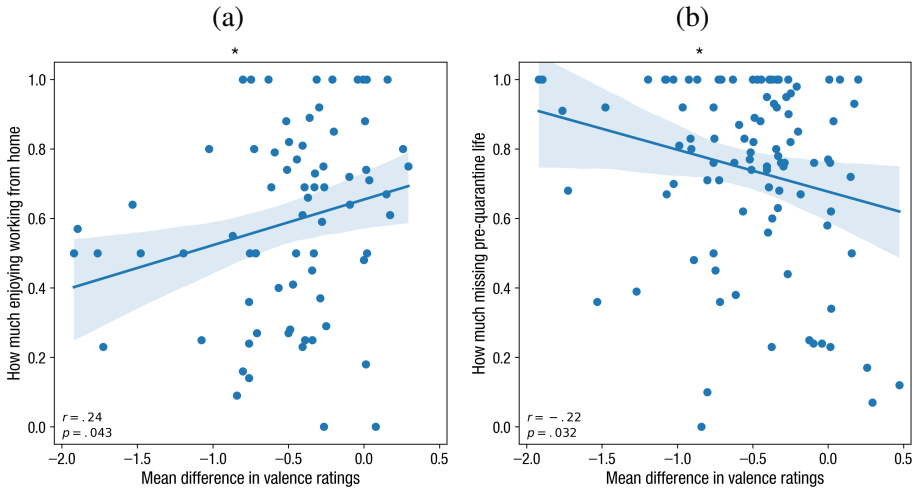


Figure 7.5: Correlations between the differences in valence ratings per participant and self-reported situation during the quarantine period. A: Linear regression between differences in valence ratings and the degree of enjoyment to work from home. B: Linear regression between differences in valence ratings and the degree of missing the “normal” pre-quarantine life. In both graphs, blue lines represent a linear regression fit; $*p < .05$

We also report a difference in the ratings of valence between those subjects who lived alone and those who lived with their families, partners, or friends ($t(98) = -2.42, p = .017, d = .611$). Specifically, we found that participants living alone rated the images significantly more negatively (Fig 7.6).

Fourth, we analyzed the time that participants took to rate each image. To do this, we computed the median rating time for each participant. The D’Agostino-Pearson normality test revealed that the rating times were not normally distributed ($p < 0.001$). Hence, similar to other studies (Whelan, 2008), we applied nonparametric statistics for the subsequent




Figure 7.6: Differences in valence ratings relative to those of OASIS between participants who, during the quarantine, lived alone and those who did not. $*p < .05$

analyses of rating times. We found a significant positive correlation between ratings times and both arousal ($r(98) = .32, p = .001$) and valence ($r(98) = 0.25, p = .012$).

Finally, we applied machine learning techniques to showcase the potential of automatically detecting users who might be at risk of developing mood disorders based on their ratings. To achieve this, we trained an SVC classifier with the valence rating information and the questionnaire's key answers. The proposed method was able to classify between those participants who lived alone and those who lived with other people with a mean accuracy of 84 % ($SD = 4$). Additionally, another SVC classifier could determine whether participants missed the pre-quarantine life with an accuracy of 65 % ($SD = 4.5$).

7.4 Discussion

In this study, we aimed at assessing the effects of the COVID-19 quarantine on the emotional state of the affected individuals. We predicted that the quarantine restrictions and, in particular, the lock-down might negatively impact mental health. It has been shown that mood deviations are reflected in the perception of affective stimuli. Hence, to test our hypothesis, we devised an online study whereby volunteers evaluated the arousal and valence of a set of standardized stimuli and compared the acquired scores with those from the original data set. We predicted that the current ratings of valence might be lower than those of OASIS, possibly due to the recruited participants' personal and social situation during the confinement.

Our results revealed that individuals who, during the experiment, were undergoing the quarantine due to COVID-19 rated neutral stimuli as significantly less pleasant when compared to the subjects who evaluated the same images during a non-quarantine period. We propose that the reported shifts in the valence ratings might be further indicative of a more general negative affective state caused by the quarantine. Indeed, we find evidence about negative changes in perception, as measured through self-reported valence ratings of visual stimuli in people with depression compared to healthy controls (Dunn et al., 2004).

Based on the acquired data, we further observed a significant effect of some of the critical aspects of our sample's personal and working situation during the self-isolation period on the reported ratings. Our results revealed a positive relationship between how much the subjects enjoyed working from home during confinement and the affective ratings. On the one hand, this finding is consistent with the literature, which demonstrates that unemployed people tend to report higher episodic sadness levels than employed people (Krueger and Mueller, 2012). On the other hand, this result might indirectly represent the effect of a decreased in-person social interaction that many jobs entail, provided that social interaction positively impacts psychological well-being (Umberson and Karas Montez, 2010).

The experience of missing regular life before the quarantine also yielded a significant effect on the negativity of the emotive ratings. We found that those participants who missed it more also experienced more substantial negative shifts in the affective assessments of the stimuli than those who missed it less. As previously demonstrated (Miles, 2015; Brooks et al., 2020; Hossain et al., 2020), we speculate that this relationship might be directly indicative of the lowered mood stemming from the negative perception of the current situation and the desire for the social distancing measures and self-quarantine to terminate. This, in turn, may be related to an increased need for both social interaction and freedom.

Furthermore, our results revealed that the ratings of valence differed depending on the participants' social living situation. Specifically, those individuals who lived alone provided more negative ratings than those living with other people. This might suggest that increased social isolation and reduced social interaction in individuals who undergo the quarantine while living alone more negatively impact their perception and, possibly, mood. Indeed, ample scientific evidence demonstrates that social isolation can result in lowered mood and depression and induce many other adverse effects on health (Hawkey and Capitanio, 2015). These effects can range from mental disorders such as depression or anxiety (Santini et al., 2020; Cacioppo et al., 2010, 2006) to cardiovascular diseases (Caspi et al., 2006; Valtorta et al., 2016). Moreover, loneliness can have detrimental effects on health through several mechanisms, including health behaviors, cardiovascular activation, cortisol levels, and sleep (Cacioppo et al., 2002). Although social isolation and loneliness are prevalent in a large proportion of the general population, affecting both younger (Matthews et al., 2016) and older (Cornwell and Waite, 2009; Shankar et al., 2011) adults, these conditions can be exacerbated or become even more strict under exceptional circumstances that force a decrease in social contact. In the case of the COVID-19 pandemic, several studies also point out a significant psychological impact, including symptoms that correspond to those found in social isolation (Wang et al., 2020b,a; Liu et al., 2020).

The above-discussed findings converge to suggest that the mitigation strategies employed to prevent the spread of the COVID-19 pandemic

are negatively impacting the emotional state of the affected individuals, which is reflected by negative shifts in the ratings of the affective stimuli. Furthermore, this pernicious effect is exacerbated by personal circumstances related to working conditions and social isolation, which, in the long term, might result in an increased prevalence of mental health conditions such as depression or post-traumatic stress disorder (Holt-Lunstad, 2017). Importantly, in the present paradigm, we focused primarily on the evaluation of neutral and positive stimuli. According to literature (Dunn et al., 2004), however, one could expect that quarantine-induced disorders of mood might also result in shifts in the negative stimuli—the hypothesis we are currently addressing in a follow-up study.

It is worth noting that our data presented variability in the relationships between the mean difference in valence ratings and both the enjoyment of working from home and the feeling of missing life from before the quarantine. This may be explained by the interaction of additional factors that were not captured by the present experiment but might have impacted the participants' emotional state. For example, personality traits might play an essential role in the ways individual participants are affected by social isolation and how they cope with it (Taylor et al., 1969; Kong et al., 2014; Zelenski et al., 2013). Furthermore, the intensity of the enforced quarantine measures was not the same for all participants, resulting variation in self-isolation. Future studies should address these limitations by controlling for additional, possibly confounding factors. Moreover, the participant sample used in this study comes from a variety of European countries. This sampling approach was intentionally chosen to cover a set of regions with comparable cultures as well as quarantine and self-isolation measures. It is possible, however, that the underlying diversity of the sample could have introduced heterogeneity in the data, which could impact the generalizability of our findings. This limitation shall be addressed in future studies by focusing the collection of data from a smaller subset of countries to further ensure the commonality of demographic aspects that could better represent the mental health of the sampled population.

On the one hand, the outcome of this study highlights the impact of

the COVID-19-induced quarantine on the affective states, thus emphasizing the need for continuous monitoring of the psychological health and well-being of the general population. Since the psychological effects of isolation might have long-term consequences, the identification of individuals at risk and carrying out interventions to mitigate the reported negative impact might be necessary not only during but also post-quarantine. On the other hand, the hereby proposed method for diagnosing the affective changes through subjective ratings of emotive stimuli may already be of use to the healthcare system. Specifically, the current findings, as well as the reported machine learning techniques, could be translated into clinical practice by using techniques such as in-person visits and digital technology in the form of smartphone apps. The former could provide a unique opportunity of combining multidimensional scales including, for instance, brain scanning (e.g., functional Magnetic Resonance Imaging) genomic measurements, observer-rated neurocognitive evaluations (e.g., HDRS), patient self-reports (e.g., BDI), medical record reviews, as well as implicit measures such as the affective evaluations used in our study. From the academic and medical perspectives, such a compound diagnosis could contribute to fundamental advances in understanding neuropsychological conditions. However, there is a need for easy to apply and low-cost solutions for diagnostics, monitoring, and treatment. Hence, the implicit assessment validated in our study can allow continuous monitoring of the effective ratings as the proxy of the affective states allowing for a prediction of the personal situation based on the obtained ratings. Such software could promote at-home remote diagnostics and monitoring of at-risk patients continuously, at a low cost, and with a further benefit of preventing possible response biases (Paulhus, 2002; Braun et al., 2001; Paulhus, 2017). We have successfully deployed such an approach in the domain of stroke rehabilitation. We have successfully deployed such an approach in the domain of stroke rehabilitation (Ballester et al., 2015; Grechuta et al., 2020). To this end, in future studies, we shall more systematically investigate the specific factors that may influence the participants' affective ratings, including personality type, as well as other symptoms that might indicate abnormal psychological states, such as insomnia. Moreover, we

will further validate the statistical relationship between the proposed implicit measure of the affective states and standard tools used to evaluate the mood, such as BDI (Beck et al., 1996) or PHQ-9 (Kroenke et al., 2001).

The efficient diagnosis, monitoring, and treatment of a neuropsychiatric illness are becoming increasingly important because its burden exceeds that of cardiovascular disease and cancer (Vigo et al., 2016) and it is estimated that about 25% of individuals will suffer neurological or mental disorders at some point in their lives. However, due to several factors, including the lack of trained healthcare professionals, pervasive underdiagnosis, and stigma, only 0.2% will be able to receive the necessary treatment (Sayers, 2001). Hence, key current challenges include the improvement of the efficacy of the diagnosis of psychological disturbances and overcoming known limitations of current clinical scales (Bagby et al., 2004; Zimmerman et al., 2005; Gibbons et al., 1993; Bech et al., 1984; Maier et al., 1985; Fried and Nesse, 2015) together with accurately capturing symptoms and patient specific concerns (Demyttenaere et al., 2020). To this end, we propose that an optimal evaluation strategy may comprise explicit, observer-rated and self-reported evaluation tools combined with implicit physiological and behavioral monitoring using biometric sensing, such as the proposed affective rating methods and associated tools (Reinertsen and Clifford, 2018).

Importantly, at the current stage, the proposed classification algorithms serve rather as proof of the potential to automatically classify well-being (Lipton et al., 2014). Future work will address this limitation by further improving the model. Those improvements will imply additional training of the classifier and the inclusion of supplementary variables that might affect participants' mental state, such as personality traits and biometrics.

Additionally, the present findings support the notion that the results from online studies carried out during the quarantine period that rely on the assessment of affective ratings or similar, might be significantly affected. Hence, this impact should be considered in the analyses and the interpretation of the acquired results.

Taken together, the present report presents a significant and timely finding which sheds light on the current quarantine's impact beyond the experience of the individuals who undergo it. In line with other studies (Nobles et al., 2020; Holmes et al., 2020; Rajkumar, 2020; Torales et al., 2020) our results confirm that individuals undergoing current mass quarantine can experience adverse psychological effects and be at risk of anxiety, mood dysregulations, and depression, which, in the long term, may lead to post-traumatic stress disorder and affect overall well-being (Miles, 2015; Brooks et al., 2020; Hossain et al., 2020). Indeed, according to previous studies, the measures that are commonly undertaken to mitigate pandemics, including stay-at-home rules and social distancing may have drastic consequences. For instance, people can experience intense fear and anger leading to severe consequences at cognitive and behavioral levels, culminating in civil conflict and legal procedures (Miles, 2015) as well as suicide (Barbisch et al., 2015; Rubin and Wessely, 2020). In addition, the long-term impact of this change in well-being is currently not understood and deserves further study. The results presented in this report highlight the need to explore possible impacts of the COVID-19 pandemic and its effects on psychological well-being and mental health. To this aim, more studies need to be conducted to systematically investigate the interventions that may be deployed by both the healthcare system and individuals undergoing quarantine to mitigate the adverse psychological effects.

Chapter 8

IMPLICIT MEASURES IDENTIFY DIFFERENTIAL EFFECTS OF THE COVID-19 QUARANTINE ON THE PSYCHOLOGICAL HEALTH OF DIVERSE POPULATIONS

This chapter is based on:

López-Carral, H., Grechuta, K., and Verschure, P. F. (2023b). Implicit measures identify differential effects of COVID-19 on the psychological health of diverse populations. *Manuscript in preparation*

The COVID-19 pandemic has had significant adverse impacts on the psychological well-being of affected populations, including depression, stress, and post-traumatic stress disorder. This situation highlighted the need for novel remote assessment tools capitalizing on digital technologies. Our study aimed to design and validate evidence-based strategies

to predict psychological states using affective ratings and other implicitly acquired signals in a remote setting. Our results show the potential for different implicit interaction features (mouse movements, keystroke dynamics, and text sentiment analysis) to be used as indicators of mental well-being. Additionally, our results reveal the differential effects that the COVID-19 quarantine in Spring 2020 had on the psychological health of the general population based on demographic factors such as gender and age. Our findings suggest that implicit interaction features can be used to monitor mental health in a scalable and cost-effective manner, with potential applications in remote mental healthcare.

8.1 Introduction

Since December 2019, the Coronavirus Disease 2019 (COVID-19) caused by SARS-CoV-2 provoked critical social, economic, and healthcare challenges (Zhu et al., 2020). With the objective of protecting the general public's health, mitigating the transmission and evolution of the virus, and reducing contagion rates, most countries worldwide enforced strict quarantine measures. Critically, growing evidence establishes that pandemic-induced containment strategies, including the lockdown and social distancing, as well as stressors such as uncertainty, financial loss, and fear of infection, have an adverse impact on the psychological well-being of the affected populations (López-Carral et al., 2020; Miles, 2015; Brooks et al., 2020; Hossain et al., 2020). Among others, the commonly reported consequences include eating disturbances, depression, post-traumatic stress disorder (PTSD), and even suicide affecting adults and children both in the acute stage and in the long term (Bai et al., 2004; Sprang and Silman, 2013; Barbisch et al., 2015; Rubin and Wessely, 2020; Bo et al., 2020; Fernández-Aranda et al., 2020). Hence, it is of critical importance to effectively diagnose the physical and psychological conditions of the general population with a focus on those at risk, on the one hand, and monitor already diagnosed patients, on the other. However, the design of robust and scalable methods for accurate diagnosis and monitoring as well

as their implementation within different healthcare systems constitutes a major challenge.

The need to develop alternative remote assessment tools becomes even more prominent in the light of evidence demonstrating a significant decrease in the number of visits to hospitals or other healthcare settings during the COVID-19 pandemic (Roca et al., 2013) and the ubiquitous financial insecurity, including unemployment, that impedes access to healthcare services. Reduced access to different services, in turn, may have detrimental effects on both physical and mental health, significantly impacting the quality of life of individuals and their families (Kawohl and Nordt, 2020; Zhao et al., 2020).

In recent years, the usage of digital health applications has shown promise as a useful tool for addressing some of the existing mental healthcare challenges. For example, smartphone applications are being used to improve mental well-being by tracking different metrics and enabling fast interventions by professionals (Woodward et al., 2020; Luxton et al., 2011). Among these features, we would like to highlight the potential of affective ratings (López-Carral et al., 2020), keystroke dynamics (Epp et al., 2011; López-Carral et al., 2019), mouse movements (Lali et al., 2014; Schaaff et al., 2012), and text sentiment analysis (Nandwani and Verma, 2021; Mäntylä et al., 2018; Serrano-Guerrero et al., 2015).

To address the ongoing mental health challenge, in this study, we aimed to design and validate evidence-based strategies to predict psychological and physical states using implicitly acquired signals in a remote setting. Our results support our previous findings (López-Carral et al., 2020) and further show the potential for different implicit interaction features (mouse movements, keystroke dynamics, text sentiment analysis) to be used as an indicator of mental well-being. Moreover, our results identify the differential effects that the COVID-19 quarantine in Spring of 2020 had on the psychological health of the general population, based on demographics factors including gender and age.

8.2 Methods

8.2.1 Participants

Our study included 300 valid participants who completed the experiment. The mean age of the participants was 31.39 ($SD = 9.33$), with 73.67 % reporting being male. Participants were recruited online using the platform Amazon Mechanical Turk (MTurk), which allowed us to reach a diverse and representative sample of participants. The demographics of our sample were similar to those of previous studies using MTurk, which helps ensure the generalizability of our findings, including our previous study (López-Carral et al., 2020) and the set of affective images that we showed to the participants (Kurdi et al., 2017).

8.2.2 Materials

In order to collect affective ratings for the images we presented to the participants, we used the Affective Slider (AS) (Betella and Verschure, 2016). The Affective Slider is a tool that allows participants to provide continuous ratings of their emotional responses to stimuli. It is composed of two sliders: one for arousal and one for valence. Both sliders are flanked by emoticons on either side to represent their extreme values. For arousal, it goes from sleepy to wide-awake, while for valence it goes from sad to happy. One slider appears over the other, with the order being randomized. Both sliders start with the handle in the middle. An example of its usage in our study is presented in Figure 8.1.

As the experimental stimuli, we used images from the Open Affective Standardized Image Set (OASIS) (Kurdi et al., 2017). OASIS is a set of 900 images with normative affective ratings of arousal and valence, similar to the well-known International Affective Picture System (IAPS) (Lang et al., 2008) but offering more updated ratings and making the images available for their usage in online experiments free of copyright. Out of the total, we selected 46 images covering the whole range of both dimensions, by setting 6 bins across each dimension and selecting 2 images

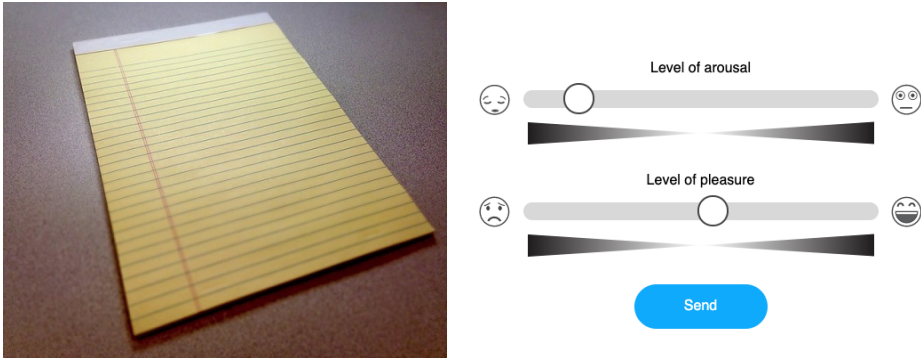


Figure 8.1: **Example of usage of the Affective Slider in our study.** The image pictured is being rated as having low arousal and neutral valence.

from each (see Figure 8.2). The same 46 images were presented to all participants, in randomized order.

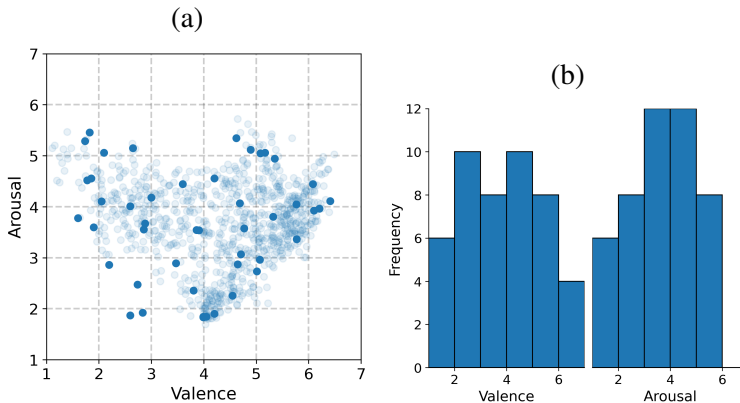


Figure 8.2: **Images selected from the Open Affective Standardized Image Set (OASIS).** A: Selected images (solid circles) in contrast with the total set (translucent circles), divided in bins over the total range. B: Histogram showing the amount of images selected for each bin, separated by valence and arousal.

To assess the potential depression severity of the participants, we used the Patient Health Questionnaire (PHQ-9) (Kroenke et al., 2001). This tool is commonly used to evaluate the symptoms of depression in clinical populations and assess their severity. The PHQ-9 consists of nine questions that are designed to measure the nine diagnostic criteria for major depressive disorder according to the DSM-5. By completing the PHQ-9, participants provided information about their potential depression severity, which we used to understand how this severity may be related to other variables, such as a bias in the affective ratings.

To assess the personality traits of the participants, we used a brief version of the Big Five Personality Inventory (BFI-10) (Rammstedt and John, 2007). The BFI-10 is a well-validated and widely-used tool for assessing the five broad dimensions of personality, known as the “Big Five” factors: openness, conscientiousness, extraversion, agreeableness, and neuroticism. The BFI-10 consists of 10 items that are designed to measure each of these five dimensions. It is a shortened version of the BFI-44, a longer inventory with 44 items (John et al., 1991). We used the personality information to understand how these traits may be related to the other metrics collected.

Finally, we used a questionnaire to assess different aspects of the participants’ experience during the COVID-19 pandemic. This questionnaire included a range of questions designed to measure different aspects of participants’ experiences, such as the impact of the pandemic on their daily lives, their mental health, and their experiences with remote work or education. Additionally, we asked participants to describe what was the best and the worst thing that had happened to them during the quarantine period. We used this questionnaire as a way to understand how the pandemic may have influenced participants’ emotional responses to the images used in the study, and to provide context for our findings. The questionnaire was completed after the other measures in the study, and participants were asked to provide their responses based on their experiences during the pandemic. Overall, the questionnaire provided valuable insights into participants’ experiences during this unprecedented time. Importantly, participants also reported whether they were undergoing quarantine. As we

wanted to focus on the impact of lockdown, we excluded from the main analyses 65 participants who reported not doing quarantine.

8.2.3 Procedure

Each participant began the study by opening the experimental website, which greeted them with a short explanation of the purpose of the study (blinding them to the actual purpose and relationship with the pandemic), instructions on how to proceed, and a brief questionnaire to collect demographic information. Then, participants were directed to the main task, which consisted of providing affective ratings using the Affective Slider for each of the 46 images selected. After completing the main task, the BFI-10 and the PHQ-9 were presented in a randomized order. Finally, the COVID-19 questionnaire was presented. Upon completion of the questionnaire, participants received an extended explanation of the experiment's rationale and were thanked for their participation.

8.2.4 Implicit Interaction Measures

While participants carried out the main task, providing the affective ratings using the Affective Slider, we tracked the position of their mouse over time. This allows us to compute the average mouse speed during this task, which has been suggested to be indicative of affective states (Lali et al., 2014; Schaaff et al., 2012). Furthermore, when participants were describing the best and the worst thing that had happened to them during the quarantine period, we tracked the keystroke dynamics of the text they typed, logging the timing of each key press and release. This allowed us to compute two relevant features from the keystroke dynamics: dwell times, defined as the time between pressing and releasing a key; and flight times, defined as the time between two successive key presses (Monrose and Rubin, 2000; Bergadano et al., 2003; Epp et al., 2011). Previous results show a negative correlation between dwell times and valence, as well as between flight times and arousal (López-Carral et al., 2019).

8.3 Results

We started the analysis by verifying that the affective ratings provided by the participants were aligned with those provided in the OASIS' normative ratings. To do this, we computed correlation analyses between the mean scores in both OASIS and our study and found significant results for both arousal ($r(46) = 0.93, p < 0.001$, Figure 8.3a) and valence ($r(46) = 0.94, p < 0.001$, Figure 8.3b).

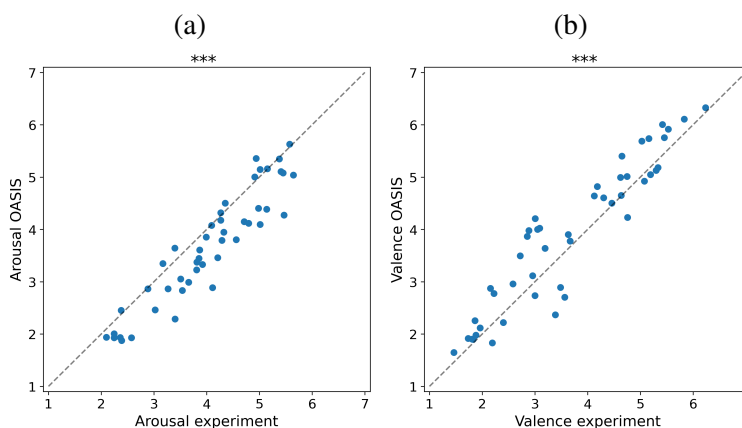


Figure 8.3: Linear correlations between the ratings obtained in our study and those from OASIS. A: Linear correlation between arousal ratings from OASIS (y-axis) and those acquired in the present study (x-axis). B: Linear correlation between valence ratings from OASIS (y-axis) and those acquired in the present study (x-axis). In both graphs, dashed lines represent the identity lines; $***p < .001$

Next, we analyzed the shift in mean ratings, comparing the affective ratings obtained in our study versus those in OASIS, for images with neutral valence (between 3 and 5, on a 1 to 7 scale). Although we found a trend towards higher arousal in our experiment, it was not statistically significant ($U(17) = 185.0, p = 0.168$, Figure 8.4a and 8.4b). However, we found a significant shift towards lower valence ratings in our experiment ($U(17) = 85.0, p = 0.042$, Figure 8.4c and 8.4d).

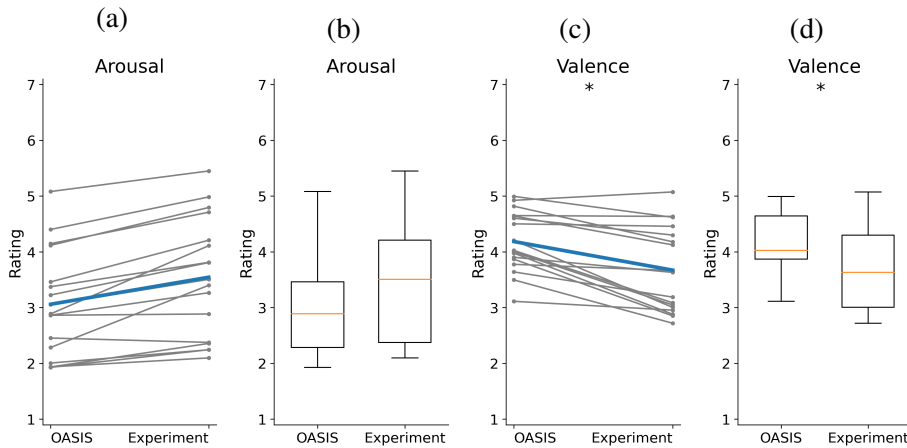


Figure 8.4: Shift in the affective ratings for neutral images. The graphs present the comparison between the ratings of arousal (A, B) and valence (C, D) for neutral images obtained in our study with those from the OASIS. In A and C, the thicker blue lines correspond to the mean, while the individual lines show differences for individual images ($N = 17$); $*p < .05$

Next, we analyzed the speed of mouse movements when providing the affective ratings using the Affective Slider, per image. We found a significant correlation between the mouse speed (measured in pixels per second) and the mean valence rating provided ($r(46) = -0.69, p < 0.001$, Figure 8.5). We did not find such a correlation with the mean arousal ratings ($r(46) = -0.14, p = 0.352$).

We then looked at the text typed by the participants describing the best and the worst things that happened to them during the then-ongoing COVID-19 quarantine period. Conducting a text sentiment analysis, we found a significantly lower mean valence when describing the worst thing in comparison to the best thing ($U(235) = 15013.0, p < 0.001$, Figure 8.6b), as expected. Additionally, we also found a significantly higher mean arousal when describing the worst thing in comparison to the best thing ($U(235) = 17224.5, p < 0.001$, Figure 8.6a). Additionally, fo-

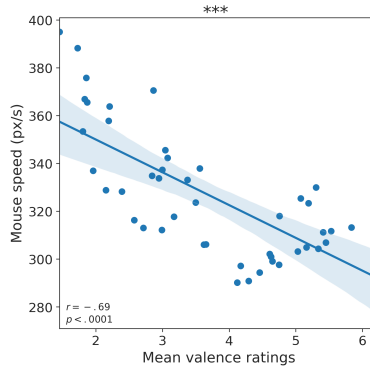


Figure 8.5: **Correlation between mouse speed and mean valence rating for each image.** * * * $p < .001$

cusing on participants with a PHQ-9 above 15 points (moderately severe depression), we found a negative correlation between the maximum text arousal and the PHQ-9 score when describing the best thing that happened to them ($r(32) = -0.44, p = 0.018$). Similarly, we found a negative correlation between the minimum text valence and the PHQ-9 score when describing the worst thing that happened to them ($r(32) = -0.35, p = 0.030$). Next, analyzing the keystroke dynamics, we found a significant difference in the dwell times between both text pieces ($U(228) = 11090.5, p = 0.049$, Figure 8.6c).

Looking at the differential features between demographic factors, we found significant effects regarding gender. Female participants had more severe depression scores than males ($U(232) = 4439.5, p = 0.020$). Moreover, female participants also provided affective ratings with lower valence ($U(232) = 4056.0, p = 0.001$), although not lower arousal ($U(232) = 5469.0, p = 0.840$). Regarding text sentiment analysis, a lower median arousal was found for female participants when describing the best thing that had happened to them during that period ($U(204) = 3335.5, p = 0.032$), although this was not the case for median valence ($U(204) = 3578.0, p = 0.133$).

Finally, looking at age differences, we also found significant differ-

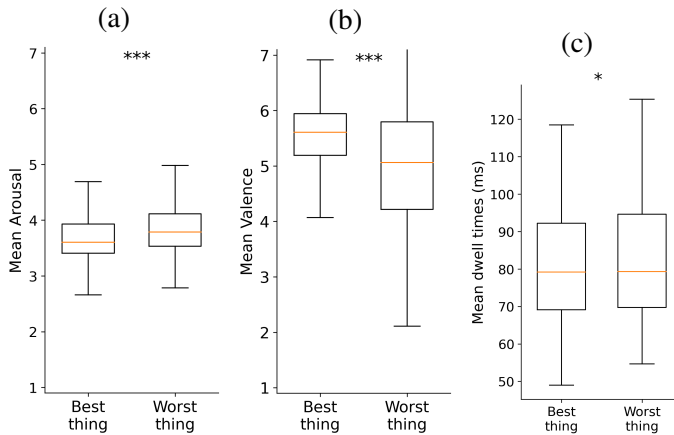


Figure 8.6: **Analysis of the text typed by the participants describing the best and the worst things that happened to them during the then-ongoing COVID-19 quarantine period.** A: Difference in mean arousal ratings between the texts describing the best and the worst thing that had happened to the participants. B: Difference in mean valence ratings between the texts describing the best and the worst thing that had happened to the participants. C: Difference in mean dwell times between the texts describing the best and the worst thing that had happened to the participants. $*p < .05$, $***p < .001$

ences regarding the age of the participants. In particular, we found a negative correlation between age and both arousal ($r(235) = -0.15, p = 0.020$) and valence ($r(235) = -0.13, p = 0.04$).

8.4 Discussion

The COVID-19 pandemic has had significant adverse impacts on the psychological well-being of affected populations, including depression and post-traumatic stress disorder. Our study aimed to design and validate evidence-based strategies to predict psychological and physical states using implicitly acquired signals in a remote setting. The results of our

study support our previous findings (López-Carral et al., 2020) and further demonstrate the potential for different implicit interaction features (mouse movements, keystroke dynamics, and text sentiment analysis) to be used as indicators of mental well-being. Additionally, our results show the differential effects that the COVID-19 quarantine in Spring 2020 had on the psychological health of the general population based on demographic factors such as gender and age.

First, we showed a significant shift in affective ratings of images from the Open Affective Standardized Image Set (OASIS) (Kurdi et al., 2017), when comparing the data we collected during the quarantine with the normative ratings from OASIS, collected before under normal conditions. In particular, there was a shift towards negativity in valence ratings of neutral images. This supports our previous results (López-Carral et al., 2020), associating a negative mental state with a negative assessment of emotional stimuli.

Next, we demonstrated the potential for different implicit features to reveal valuable insights of affective states. We showed that mouse speeds were negatively correlated with the mean valence ratings. From this, we can interpret that participants behaved faster to provide the ratings for more unpleasant images, possibly to stop visualizing that image and advance to the next one quickly. These results show promise, although a specific study focusing on mouse dynamics and controlling other factors would be needed.

Regarding text provided by the participants, we showed that we can infer emotional information using two different methods. Not only by looking at the content typed, i.e., text sentiment analysis, but also at how it is typed, i.e., keystroke dynamics. We found one feature of keystroke dynamics, dwell times, that was differentiated depending on whether the participants were describing something positive or negative. This is aligned with previous results that correlate dwell times with valence (López-Carral et al., 2019).

Finally, we inspected how the lockdown situation had affected different population groups differently. We found that female participants not only had higher severity of depression, according to their PHQ-9 score,

but also that they rated images as less arousing and more negative than male participants. This goes in line with existing studies that point towards a more significant impact of the pandemic on the mental health of women (Dal Santo et al., 2022). Regarding age, we also found a difference regarding how the images were rated. Overall, older participants also rated images as less arousing and more negative than younger participants. This also supports evidence that the pandemic affected older citizens more strongly, both physically and mentally, with poorer mental health (Wilson et al., 2021; Banerjee, 2020).

Overall, our results from this study not only offer insights into the specific psychological impact of the COVID-19 lockdown of the Spring of 2020, both generally and for specific populations, but it also proposes an array of new digital-first methods to measure and monitor mental well-being.

Digital-health solutions have been demonstrated to have the potential to effectively address current challenges faced by the healthcare system and preserve access to essential medical services (Fagherazzi et al., 2020; Sust et al., 2020). Such tools may be implemented to complement the standard practice with the advantage of providing continuous monitoring tailored to individual users and their needs. Our findings suggest that implicit interaction features can be used effectively to predict mental well-being in a remote setting, providing a valuable tool for assessing and monitoring mental health during the COVID-19 pandemic. However, further research is needed to fully understand the implications of our findings and to develop more robust and scalable methods for accurately diagnosing and monitoring mental health in a remote setting.

Chapter 9

BINARY AFFECTIVE RATINGS THROUGH SWIPING

This chapter is based on:

López-Carral, H. and Verschure, P. F. (2023). Binary affective ratings through swiping. *Manuscript in preparation*

The circumplex model of emotion, which decomposes affect into arousal and valence, is widely employed in different scientific fields studying emotion. In order to assess these, different tools exist, such as the pictorial Self-Assessment Manikin or the digital Affective Slider. While these methods work well to obtain precise ratings, they are not intuitive to all participants and require understanding the concepts of arousal and valence to provide separate appraisals of both. In order to overcome these limitations, here, we are presenting a novel method of obtaining binary affective ratings through swiping using a mobile app, called the affective swiping app (ASA). We validated our approach in an online study with 303 participants in which we asked participants to swipe right or left to indicate whether they liked or disliked a series of images from

a standardized set that includes normative affective ratings. Our results show a robust correlation between the like percentage for each image and their normative valence rating on a continuous scale. Furthermore, we found that implicit measures like response times and the swipe velocity applied diverged correlated with absolute valence polarity, diverging from the neutral value of valence. Overall, our results demonstrate that ASA provides for a rapid and intuitive method of rating stimuli could provide rich and reliable affective ratings, making it suitable for different types of studies and applications, including mental well-being monitoring.

9.1 Introduction

Emotions have been a topic of study in psychology for decades, with multiple theories and models having been proposed over the years. One of the most widely used models of affect is the circumplex model proposed by James A. Russell in 1980, which establishes affect as comprising two separate but interrelated two dimensions: arousal and pleasure (also referred to as valence) (Russell, 1980). In its original form, these two dimensions are assumed to be orthogonal, with different emotions being defined in the combination of their relative valence and arousal bipolar weight. For example, “happy” would correspond to relatively high arousal and pleasure, while “relaxed” would correspond to similarly high pleasure but low arousal. On the opposite end of the pleasure scale, “angry” is placed as high arousal but low pleasure, and “depressed” with low arousal and low pleasure.

The bipolar valence and arousal model has been widely employed, beyond psychology and into neuroscience (Posner et al., 2005). Yet, in recent years, different studies have investigated the relation between valence and arousal, questioning their independence (Kuppens et al., 2013). While different types of relations have been proposed, a recent review concluded there is a weak but consistent *V*-shaped relation of arousal as a function of valence with a large variation at the individual level that can modulate their relation (Kuppens et al., 2013). Indeed, a recent study mea-

sured brain activity using functional magnetic resonance imaging (fMRI) suggested that arousal is not separable from valence and instead might be representative of the intensity of valence (Haj-Ali et al., 2020). Databases of images with ratings for affect, such as the International Affective Picture System (IAPS) (Lang et al., 1993, 2008) or the more recent Open Affective Standardized Image Set (OASIS) (Kurdi et al., 2017) have also found this relationship, in which the highest arousal ratings are found at the positive and negative extremes of the pleasure dimension. Recent studies working with this model have resorted to asking study participants to rate only one or the other to avoid contamination between the two dimensions (Kurdi et al., 2017). Hence, the assumed orthogonality and independence of the arousal and valence dimensions of the classic bipolar model of affect is rather an hypothesis that is not fully empirically established.

A popular tool for the self-report of affect is the Self-Assessment Manikin (SAM), proposed in 1994 as a method to report pleasure, arousal, and dominance (Bradley and Lang, 1994). SAM is a non-verbal, pictorial tool that uses 5 drawings per dimension to represent different values on a discrete scale. Pleasure is represented with a character ranging from a frowning face to a smiling one, while arousal has a character with a pattern resembling an increasingly large explosion at its center, and dominance is depicted by a character that incrementally increases in size. SAM was designed for pencil-and-paper responses, with a computer version also created. In order to overcome some of SAM's limitations (including discrete ratings, potentially confusing graphical depictions, and non-digital-first nature), the Affective Slider (AS) was introduced in 2016 as a digital tool to report arousal and pleasure on a continuous scale (Betella and Verschure, 2016). The AS was proposed as a potential replacement for the SAM, providing equivalent ratings while also simplifying the rating process and making it easily reproducible on new digital devices. It consists of two separate sliders with emoticons on each side of both: unhappy/happy faces for pleasure, and sleepy/wide-awake faces for arousal. It was shown that AS provides values which are consistent with SAM and since its introduction it has gained wide following.

While the AS succeeds in enabling relatively simple and precise affective ratings on digital devices, its usage still requires the conceptual understanding of both arousal and pleasure by the users. To attempt to ensure that understanding, experiments using the AS might rely on textual or oral explanations of both dimensions as part of the experimental design. This puts an additional burden on participants to understand these concepts and then, for each stimulus, briefly reflect on them to provide two separate ratings, introducing a potential for cognitive biases in the ratings provided. In addition, such an explanatory phase adds further complexity to experimental protocols. Hence, the question is whether new approaches can be leveraged to overcome these limitations of the AS and its underlying circumplex model of affect.

In recent years, with the increasing popularization of digital media, portable devices, and touchscreens, a rating method has become increasingly widespread for a variety of use cases: swiping. In the popular online dating app Tinder, users swipe right or left to indicate whether they like or dislike (respectively) other users that are presented as potential dating candidates. Other apps also adopted this rating method, including non-romantic dating, job finding, music discovery, and clothing (Kerckhove and Pandelaere, 2018; David and Cambre, 2016). This raises the question whether this simple method of expressing binary preference can be leveraged for a more effective assessment of affective stimuli and, with it, the mental state of the users.

To provide fast and intuitive affective ratings, we are presenting a novel method of binary ratings through swiping, designed for mobile touch-screen devices. Here, we describe an online study using this methodology and its associated affective swiping app (ASA) to provide like-dislike ratings of images from a standardized affective stimuli set. Our results show that by aggregating the intrinsically binary ratings, we can obtain affective information in a continuous range approximating the ratings provided using a more elaborate and continuous scale, such as the AS. Furthermore, implicit measures related to the swiping interaction (timing and velocity) are also indicative of the normative affective polarity. Given these results, we propose the ASA method as an alternative for

rapid, mobile-first approaches to obtaining affective ratings in which simplicity, speed, and portability are the priorities. This could be of particular relevance to studies harnessing the benefits of smartphones for health (mHealth) and mental well-being (Luxton et al., 2011), facilitating quick and frequent affective assessments. Indeed, previous results suggest that affective ratings might be an indicator of disrupted mental health (López-Carral et al., 2020).

9.2 Methodology

We collected data from a total of 303 participants that successfully completed the experiment online. We recruited participants on Amazon’s Mechanical Turk (Amazon, US) platform. They received 1 USD for their participation. 57.76 % of participants were male, with a mean age of 36.22 ($SD = 10.21$). 79.21 % of them were from the United States, with the rest being from 9 other countries.

After opening the website for the experiment using the address provided, participants were greeted with a page that described the task and provided instructions. Here, participants were also asked to report their gender, age, countries of citizenship and residence, and level of education. Participants were asked to complete the experiment using only a smartphone and a technical assessment was in place to ensure participants did not use other devices. 70.63 % of participants used an Android phone, with the remaining 29.37 % using an iPhone.

After providing their consent for participating, the main task started. The task consisted of providing binary affective ratings by swiping images right (like) or left (dislike) (see Figure 9.1). A total of 100 images were consecutively presented, in a randomized order. Participants were instructed to not think about their decision for long and instead respond immediately. In addition to the image assessment trials, participants were also presented with 20 control trials that explicitly asked them to swipe in a particular direction. The first 10 were presented immediately before the image trials, while the remaining were randomly interspersed during the

rest of the task, rendering a total of 120 trials.



Figure 9.1: Affective swiping app (ASA) and the method applied to rate the images. Participants were presented with images individually in a virtual stack of cards and were asked to swipe them right if they liked them or swipe them left if they disliked them. As the participants swiped right or left, the image card got a green or red overlay, respectively, and the next image could be seen behind. The images shown belong to OASIS (Kurdi et al., 2017).

The images used were from the Open Affective Standardized Image Set (OASIS) (Kurdi et al., 2017), which provides a collection of images with normative ratings for arousal and valence on a scale from 1 to 7. This set is similar to the well-known International Affective Picture System (IAPS) (Lang et al., 2008), but features more recent images and ratings and can be freely used online without copyright restrictions. Out of the 900 images in OASIS, we chose a subset of 100 images by binning the set into 6 by 6 equal bins along both dimensions (arousal and valence, see Figure 9.2a) and randomly selecting images based on the valence dimension, with 60 % being neutral (valence between 3 and 5) and

the remaining either negative (valence between 1 and 3) or positive (valence between 5 and 7) (see Figure 9.2b). This ensured a balanced subset between both halves of the valence dimension, with an emphasis on neutral images. The resulting subset of images, like the full OASIS (Kurdi et al., 2017), shows a V-shape relation between arousal and valence (see Figure 9.2c, valence lower than 4: $r(50) = -0.44, p = 0.001$, valence higher than 4: $r(50) = 0.53, p < 0.001$). This is representative of the potential underlying relationship between both dimensions, as mentioned earlier.

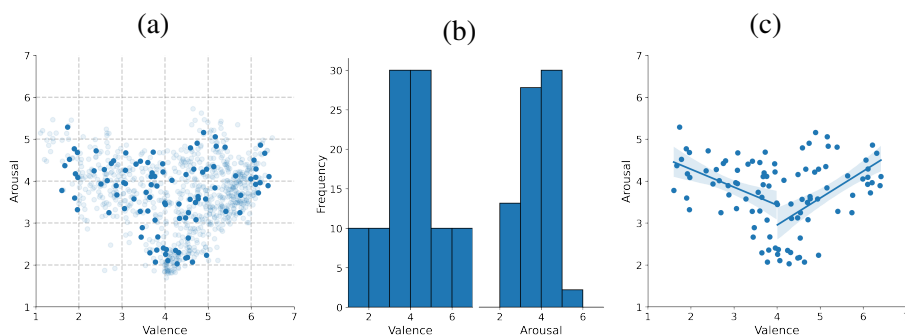


Figure 9.2: Subset of images selected from OASIS (Kurdi et al., 2017). (a) The distribution of 100 images used in our study (solid circles) from the full set of 900 (semitransparent circles) in the valence and arousal space. (b) Histogram of selected images across the valence and arousal dimensions. The sampling was performed based on the valence ratings. (c) Relation between arousal and valence in the subset of selected images, showing a replication of the overall asymmetric V-shape structure found in OASIS, i.e., a negative correlation for above-average valence scores and a positive correlation for below-average scores.

For each swiping interaction, we computed the median response time, measured as the total time elapsed in each trial (from the image presentation until the completion of the response, when the finger is released from the screen), and the median swiping speed during the touch gesture, measured as the movement speed of the touch gesture in pixels per second.

After the main task was completed, participants were also asked to answer a brief depression questionnaire, PHQ-9 (Kroenke et al., 2001), in order to assess the potential severity of depression across the sample. After this was completed, participants were thanked for their participation and debriefed with an explanation of the experimental rationale, concluding the experiment.

9.3 Results

First, we analyzed the responses that were provided by the participants by swiping the images. In particular, we computed the percentage of left versus right responses for each image. We found a very strong correlation between the valence ratings reported in OASIS and the percentage of like responses ($r(100) = 0.90, p < 0.001$) (see Figure 9.3a). This correlation was high with a random subset of participants, and rapidly increased further as more participants were included in the analysis (see Figure 9.3b). A k -nearest neighbors model could predict the swipe response based on valence with a 74 % accuracy for the full set of images, 88 % accuracy for positive images (valence over 5) and 85 % accuracy for negative images (valence under 3).

We analyzed the implicit metrics against the absolute polarity of valence, i.e., the difference from the central value (4, in the original scale from 1 to 7). We found a significant negative correlation between the absolute polarity of valence and response time ($r(100) = -0.46, p < 0.001$, Figure 9.4b), with a positive correlation for negative valence ($r(50) = 0.41, p = 0.003$) and a negative correlation for positive valence ($r(50) = -0.51, p < 0.001$) (see Figure 9.4a). Therefore, participants took more time to respond to more neutral images, while they responded quicker to more extreme values of valence. The average response time overall was 0.85 seconds ($SD = 0.13$). Furthermore, we found a similar effect on the swipe velocity applied to the image cards, with a significant correlation between this metric and the absolute polarity of valence ($r(100) = 0.33, p < 0.001$, Figure 9.4d), once again with varying strength

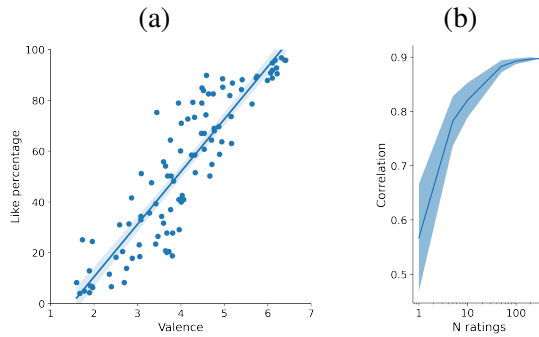


Figure 9.3: [Relationship between the percentage of like choices and continuous valence. (a) Correlation between the like percentage per image and the valence ratings from OASIS. (b) Increase in correlation between like percentage and valence as a function of the number of observations.

for the negative side of valence ($r(50) = -0.54, p < 0.001$) and the positive one ($r(50) = 0.23, p = 0.104$).

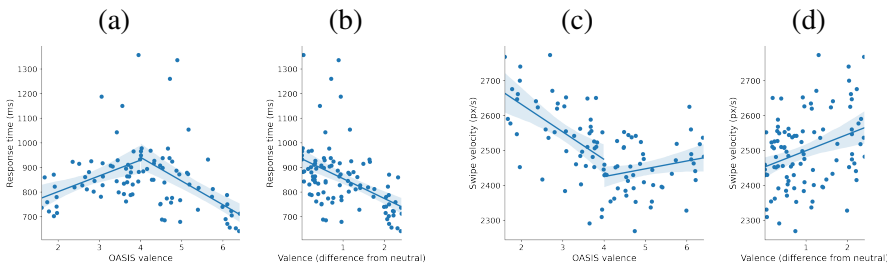


Figure 9.4: Implicit measures correlate with absolute valence polarity. (a) Split correlation between the median response times per image and their valence rating reported in OASIS. (b) Correlation between the median response times per image and their valence difference from neutral (central value of 4 in the scale). (c) Split correlation between the median swipe velocity applied per image and their valence rating reported in OASIS. (d) Correlation between the median swipe velocity applied per image and their valence difference from neutral (central value of 4 in the scale).

No significant correlations were found between the previous features and the arousal ratings provided by OASIS. Additionally, no significance correlation was found between these features and the PHQ-9 scores of the participants.

9.4 Discussion

The circumplex model of emotion, which decomposes affect into the orthogonal dimensions of arousal and valence, is widely used across different scientific domains (Kuppens et al., 2017). Although this model posits arousal and pleasure as two separate, orthogonal dimensions, more recent research has identified a clear relationship between them, with arousal being maximum at the extreme values of pleasure (Kuppens et al., 2013).

The assumption behind the circumplex model of bipolar structure of emotion has also influenced the tools used to assess affective state. Among the different tools for the self-report of affect that exist, one of the most popular is the Self-Assessment Manikin (SAM), dating from 1994 (Bradley and Lang, 1994). A more digital-focused tool was proposed in 2016, the Affective Slider (AS), consisting of two continuous and independent scales for pleasure and arousal (Betella and Verschure, 2016).

As an alternative to existing methods and their potential drawbacks, we have proposed a new method for intuitive and fast affective ratings of images through single binary swiping interactions: the affective swipe app (ASA). In this method, participants swipe right to indicate they like the stimulus and left to indicate they dislike it. Our results show that by using implicit interaction data and the proportion of positive versus negative responses, accurate affective information can be collected. In particular, we found that the mean like percentage for each image using this method correlates strongly with the valence ratings provided on a continuous scale on the standardized set of images that we employed, derived from OASIS (Kurdi et al., 2017). Furthermore, the median response time per image was indicative of the absolute valence polarity of the stimulus. We also observed that participants rated images with ex-

treme valence quickly, while they were slower for images with neutral valence. This seems to suggest that more effort had to be placed on assessing content that did not elicit a strong aversive or pleasant response. In line with this, the median velocity applied to the image cards to provide the rating was also indicative of the absolute valence polarity, with stronger responses to content with more extreme values of valence.

Interestingly, we did not find a significant direct correlation between our explicit and implicit measures and the arousal ratings reported in OASIS. Yet, our results for the implicit measures follow a relationship with valence, with *V*-shape and inverted *V*-shape trends that are analogous to the OASIS reference dataset. Similarly to arousal scores of the image set employed, response times and swipe velocities diverge as valence becomes more extreme towards either side. This questions the orthogonality of arousal and valence originally proposed in the circumplex model, and supports other studies that have questioned this assumption (Kuppens et al., 2013; Haj-Ali et al., 2020).

Additionally, a potential artifact of more traditional methods for the self-assessment of affect is that they might rely on metacognitive processes, requiring people to reflect on the affective content of the stimuli based on the concepts of arousal and valence. In contrast, the method proposed here is more immediate and thus be biased toward immediate implicit emotional appraisal. Indeed, a recent study using functional magnetic resonance imaging (fMRI) found that representing the affective value of a reward on a continuous scale might invoke separate processes as opposed to those involved in making a binary decision about whether to choose that reward, involving different areas of the prefrontal cortex (Grabenhorst et al., 2008).

Our results show that our method for binary affective ratings through swiping has the potential to be a useful tool in a variety of studies that rely on emotional self-assessment. The rapid and intuitive nature of this method makes it especially suitable for mobile-based research in which participants have to assess images repeatedly and without the need to understand concepts such as arousal and valence. An example of this would be a mental health monitoring app, in which the state of the par-

ticipants would be longitudinally monitored by their affective ratings of images. In particular, their ratings, especially for neutral stimuli, could be a reliable indicator of depressive states (Sloan et al., 1997, 2002; Dunn et al., 2004). Previous research suggests that lower valence ratings of neutral images could be indicative of worsened mental health (López-Carral et al., 2020). Although we did not find this relationship between the ratings obtained here and the PHQ-9 score of the participants, it is possible that a longitudinal study with repeated participation, other questionnaires (such as one focused on mood), or a clinical group or participants, would provide better indicators of emotional alterations. Future studies might also consider additional factors, such as the handedness of the participants (Casasanto, 2009), the swiping direction for approach-avoidance (Cervera-Torres et al., 2021), or additional implicit factors, such as keystroke dynamics (López-Carral et al., 2019).

Part III

Interactive and Adaptive Systems

Chapter 10

AN ASSISTIVE SYSTEM FOR TRANSFERRING DOMAIN KNOWLEDGE TO NOVICE USERS

This chapter is based on:

López-Carral, H. and Verschure, P. F. (2022). An Assistive System for Transferring Domain Knowledge to Novice Officers. *European Law Enforcement Research Bulletin*, 22(6)

Instructional strategies in many operative fields have reached a high level of complexity due to dynamically changing task environments and the introduction of different technologies to help users in their operational work. In the last decades, a transition has been observed from dedicated trainers to the adoption of automated technologies to support the trainees. Based on a review of state-of-the-art literature and direct feedback from the final users, we have developed an assistive system to aid in the knowledge transfer from expert to novice users and, consequently, improve the time necessary to train new practitioners. This system is grounded on the

most relevant instructional principles derived from cognitive and learning theories. The result is a system that can dynamically deliver suggestions based on previous successful actions from other users and the current performance and state of the user. To validate our system, we implemented a knowledge graph exploration task, in the context of a law enforcement task. The novel knowledge transfer system is introduced here by presenting the results from our literature review, explaining the architecture of the assistive system, and discussing our observations from the validation task. With this work, we aim to facilitate the transfer of domain knowledge, which could have a significant impact on the training and education of new users of data exploration systems.

10.1 Introduction

Instructional strategies in many operative fields have reached a high level of complexity due to dynamically changing task environments and the introduction of different technologies to help users in their operational work. In the last decades, a transition has been observed from dedicated trainers to the adoption of automated technologies to support the trainees. This paradigm shift makes transferring precise knowledge to novice users a challenging problem, which becomes especially relevant when the user is dealing with large and complex datasets from which to extract relevant information.

Supportive technologies, such as recommendation systems, have attracted a lot of interest in the last decades, both in industry and academia. The goal of such systems is to help the users to reduce the burden imposed by the high information load that is intrinsic to the exploration of large and complex datasets by providing valuable suggestions in the form of specific items or possible actions to choose from. Despite clear technical advances witnessed in the field in improving the accuracy of the recommendations, several challenges and open issues remain, especially regarding the specific role of various human factors.

This study was framed within a project aimed at providing Law En-

forcement Agencies (LEAs) with a set of automated tools and systems to boost the investigative work in the fight against illicit trafficking activities. One of the necessary functionalities identified was the capability to provide adequate solutions facilitating the transfer of the acquired expertise among experienced users and, consequently, boosting the take-up time necessary to train new users. In order to accomplish this task, we decided to build a novel assistive system, which, combining practical knowledge from classical recommender systems with theoretical knowledge from cognitive systems, is able to aid in the transfer of domain knowledge to novice users.

We will discuss, firstly, the recommender systems in general before outlining the recommender system for assisting knowledge transfer that reflects the best practices, approaches, and directions in the respective law enforcement domain. Our recommender system is conceptually grounded in a cognitive architecture, learning from interactions to later assist novice users by suggesting key pieces of information that other users have selected. Then, we describe the case used for validating this system in a knowledge graph exploration task based on a novel interface for LEAs to present the collected information in a criminal investigation. Finally, we will put forward our conclusions and outline possible next steps.

10.1.1 Introduction to recommender systems

Recommender systems have been used extensively in research and industry since the mid-1990s (Goldberg et al., 1992). The most common domain for their use is electronic commerce (e-commerce), the entertainment and media industry, and services. Many online businesses employ dedicated algorithms to provide recommendations to their customers based on inputs such as their history of items visualized and purchased or their demographic data. Another popular area in which recommender systems are used is multimedia applications (Ge and Persia, 2017). For example, many online music platforms use them to recommend songs or artists based on what each individual listens to (Song et al., 2012). Similarly, recommendation systems are common in online video platforms to

provide personalized suggestions for TV shows, movies, and other videos (Asabere, 2012).

Several types of recommender systems have been proposed that, depending on the techniques employed, can be classified into different categories (Park et al., 2012; Villegas et al., 2018; Ricci et al., 2011; Adomavicius et al., 2011). In content-based recommendation schemes, the system learns to propose items similar to those that were preferred in the past by the same user. In contrast, collaborative filtering approaches recommend items that other users with similar profiles have preferred in the past. Knowledge-based systems recommend options based on specific domain knowledge about how certain features meet users' needs and preferences. Finally, hybrid systems are based on the combination of the techniques mentioned above to improve performance (Burke, 2002).

Despite providing varying degrees of support, overall, recommender systems are not always tailored to specific user needs and situations. It has been suggested that adaptive recommender systems should be modeled in terms of situations rather than knowledge structures (Adomavicius and Tuzhilin, 2005; Richthammer and Pernul, 2020; Adomavicius et al., 2011). Such a system would be capable of delivering better results to the user by taking into account contextual factors in the delivery of highly tailored information. Typically, these contextual factors include location, time, computing context, the activity of the user, or social relations (Verbert et al., 2012).

However, context can also refer to the motivational, cognitive, and emotional aspects that are inherent to the interaction between the user and the system. Most of the research on personalized recommender systems has been focused mainly on technical issues, neglecting the importance of the underlying psychological and implicit factors when exploring and analyzing data (Buder and Schwind, 2012).

Thus, it is now considered relevant that for a recommender system to be effective, it should merge a variety of techniques and features in order to offer valuable support and reduce the demands imposed by information load. In this sense, systems have been developed that incorporate adaptive content presentation and adaptive navigation support (Brusilovsky, 2007).

Content adaptation adjusts the presentation of the content to the user's goal, knowledge, and other information, which is stored in a model of the user to balance factors such as cognitive load, arousal, or learning style (Jin et al., 2017).

10.1.2 Recommender System for Domain Knowledge Transfer

This literature review on knowledge transfer systems reveals a multifaceted and active field where a plethora of technological approaches have been proposed and developed. It also becomes apparent that individual differences (such as motivational and emotional ones) have not received proper consideration when defining effective recommender technologies. This is mainly because of a lack of coherent principles derived from learning and cognitive sciences to guide the development of such systems.

Instead of working from a pure computer science perspective, the proposed recommender system will be grounded on cognitive theories, specifically, the Distribute Adaptive Control (DAC) theory of mind and brain (Verschure, 2012). This theory will serve a dual role in the theoretical framing and the implementation of the core functionalities of the system.

DAC considers humans themselves as adaptive systems that react and adapt to the changing demands of the environment by applying self-regulation strategies in response to intrinsic goals and motivations. The same principles play a foundational role in the implementation of more effective cognitive artificial systems.

Conceptually, this recommender system can be realized as an artificial agent whose reasoning and memory components need to extract relevant knowledge from sequences of interactions in a coordinated way. The proposed system thus emerges as the interplay of the Reactive, Adaptive, and Contextual layers as defined in the DAC architecture (see Figure 10.1).

The recommender system emerges as the interplay between the three layers, which work at different timescales, with the fastest layer at the bottom and the slowest one at the top. In this architecture, the layer at the

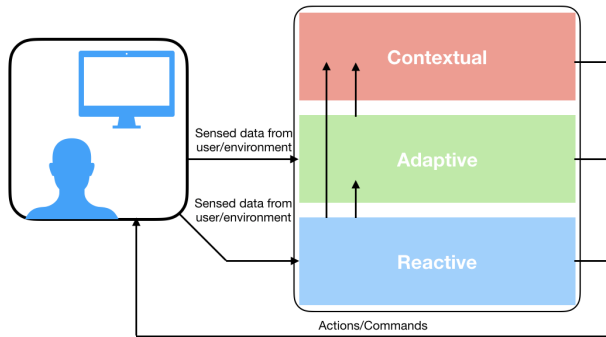


Figure 10.1: Abstract conceptualization of the cognitive architecture of the knowledge transfer system based on the DAC framework.

bottom (Reactive layer) provides the basic form of interaction, taking as input information from the environment and the user to facilitate the basic interaction.

Secondly, the Adaptive layer oversees adjusting the information given to the user, such as suggesting a specific piece of information or directing attention to a specific subset of information. Finally, the Contextual layer operates on longer timescales to learn from interactions from all the users, building profiles and detecting interaction strategies in order to create a knowledge base on which to optimize its behavior to improve its capabilities in assisting the users. All in all, the system works hierarchically at different time scales, from the immediately reactive, to the medium to adapt to each user, to the long one across different interactions.

Next, one of the key aspects of a recommender system like this, which participates dynamically during an interactive task, is to decide when to provide a suggestion. There are many criteria that could be employed, depending on factors such as the specific task that the user is carrying out, how the interface has been implemented, or the number of user feedback sources available. Although we could include more complex features related to the user state (e.g., estimating stress, attention), here we present the interaction features that we have implemented in the current version

of the system, to be used in an online task running on a web browser.

One of the interaction criteria is based on time. If the user has spent more than a specific amount of time without interacting with the system (by clicking somewhere), a suggestion is provided. This is done to stimulate interaction with the system, which is based on exploration to obtain information. This time threshold was fixed at 10 seconds.

Another criterion to provide a suggestion is based on the number of clicks that the user has performed without advancing in the given task. If the user has clicked a certain number of elements without getting closer to solving the task, a suggestion is provided with the goal of reorienting the user toward more relevant information.

If these criteria are not met, no suggestions are provided, as this would indicate that the user is carrying out the task successfully: with fluidity and accessing information that is relevant to solve the task at hand. This way, expert users, who already have successful strategies to accomplish the task, are not encumbered by unneeded recommendations, while novice users, who have not yet developed successful strategies, get the necessary guidance.

Another important aspect of the recommendation system is that not all the suggestions are equally revealing of the next action to take. Instead, there are different levels of recommendations, which are adapted dynamically based on the performance of the user. First, the system starts by providing general recommendations based on the content that just some users interacted with, but not most of them. As the users keep interacting with the system, if they have already received several suggestions at the current level, the recommendation level gets upgraded, and, consequently, the system recommends content of increasing popularity among the previous expert users who successfully solved the task.

To bootstrap the recommender system, some initial interaction data was needed. To achieve this, a custom synthetic data generator was developed. For a given task, the algorithm that was developed generates a random solution resembling one that an expert user would perform. This synthetic interaction data arrives at a solution by following a series of steps that are close to the optimal ones, by following some natural strate-

gies that most users would develop after familiarizing themselves enough with the system (i.e., becoming expert users).

The algorithm creates this synthetic data by working backward from the solution of the task (i.e., starting at the end of the interaction). Then, it generates data corresponding to clicks of random pieces of information at different levels of separation from the solution. The result is a data file almost indistinguishable from the one obtained from actual interaction data.

Finally, the last step in the process is generating the recommendations from the interaction data collected or generated. To achieve this, a custom algorithm was implemented. It gathers all the existing interaction data for a given task and lists all the existing pieces of information. Then, it counts how many times each piece of information was selected by the users. The result is a data file that will later be processed by the main application to create a ranking of possible recommendations based on this information.

10.2 Use Case: Investigation Knowledge Graph Exploration

As the initial use case of this recommender system, we chose the exploration of different knowledge graphs. These knowledge graphs represent, conceptually, one investigation. Each knowledge graph is composed of a number of interconnected nodes. The nodes represent a piece of evidence that is related to others. This is indicated by lines (edges) connecting the nodes bidirectionally.

Thus, a knowledge graph here is an abstract graphical representation of all the information collected in an investigation. This modality of information presentation and exploration was designed in collaboration with Law Enforcement Agencies as part of a bigger system of state-of-the-art tools to assist officers in their investigative work by exploiting the latest digital technologies.

In this context, to validate the resulting recommender system that we implemented, we developed a simplified knowledge graph tool that does

not use real investigation data, but a gamified and goal-oriented version of crime investigations. The users are asked to put themselves in the position of an investigator who must solve a series of investigations using a new visual interface. For this, they are invited to interact with the knowledge graphs, interacting with the nodes (again, each representing a piece of evidence) in search of a target node. This target node is the solution for each of the cases, representing the piece of evidence required to solve the investigation. Nodes around this target provide hints that allow participants to find out the solution.

Although this task uses the analogy of solving a case, it is important to emphasize that this is just the conceptual idea. As explained, the task is a simplified version, being closer to a game than an actual job of an officer investigating a real case. The way to solve each of the tasks is based on solving a series of logic puzzles, as explained below.

Figure 10.1 depicts the user interface that we implemented to present the task. The knowledge graph itself occupies the central part of the screen. Users can interact with the graph by clicking on the different nodes to obtain information about them (name and possible relationship to the target node). The name of the node also appears when hovering the mouse cursor over it. It is also possible to move the nodes by clicking and dragging, which might be helpful to get more clarity on the connection to other nodes, although this is never required. Users can also displace the graph by clicking and dragging on an empty space, as well as zooming in and out by using the controls provided in the top-right corner, although these actions are not required either. Finally, in the top center, the category of the target node is displayed.

We decided to use four different categories of nodes to provide enough diversity without being too distracting or overloading. These four categories are: person, vehicle, text, and location. Each category is differentiated from the others by using a different iconic figure and color (see Figure 10.2 and Figure 10.3).

As mentioned before, the nodes surrounding the target provide relevant information that is needed to the solving the case. Depending on their closeness and relevance, four levels are established and displayed in




Figure 10.2: The user interface of the knowledge graph exploration task. Interaction controls for the displacement and zoom of the graph are located in the top right corner (from the bottom up: zoom out, zoom in, restore view). On the top center, the category of the current target is indicated. The panel on the top left provides information about the node that is currently selected, which appears with a black outline and black connections in the graph. This panel also has the button to submit the solution, corresponding to the node currently selected.

the node information:

- “This [category name] is suspicious”: This appears for the target node and for all nodes of the same category that are within three degrees of separation from it.
- “This [category name] is directly related to the target”: This appears for nodes that are directly related to the target (first-degree connection), of a different category from it.

- “This [category name] is indirectly related to the target”: This appears for nodes that are indirectly related to the target (second-degree connection, which is, connected through exactly one node in between), of a different category from it.
- “Unclear”: This appears for all other nodes not covered in the previous three categories, i.e., all nodes that are too distant and unrelated from the target.

Using this information provided by the different nodes close to the target, the solution is implied. In each knowledge graph, there is only one possible solution, and the information, when enough nodes are explored, points unambiguously to it. Users must integrate this information in a logical manner. It is a matter of logically inferring the solution by integrating some simple relationship data.

The complexity of the task is modulated by the size of the knowledge graph, determined by the number of nodes and connections. The higher the number of nodes and connections, the higher the difficulty, as the visual complexity increases and there are more nodes to explore. We created three difficulty levels according to this: 50, 100, and 200 nodes and connections. Two graphs of each difficulty level are presented, in increasing difficulty, for a total of six cases for each participant.

As indicated, one of the key aspects of this system is the presentation of suggestions to the users. These suggestions are provided in the form of recommended nodes based on the actions of other users. When a node is suggested, it gets selected with a thicker light-blue outline. Its connections to other nodes also appear in the same color (see Figure 10.3). When a node is suggested, a panel appears on the bottom of the screen, alerting users of this fact and thus ensuring that they notice the suggestion. This message stays on the screen for 3 seconds.

10.3 Discussion and Conclusion

Here, we have presented a novel assistive system capable of learning from interactions with users in order to provide relevant suggestions to other

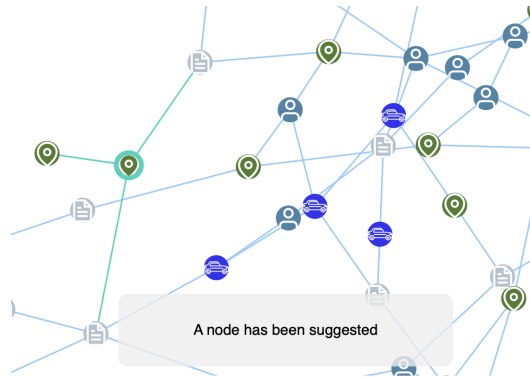


Figure 10.3: Example of a suggested node. It appears with a thicker light blue outline, as well as its direct connections to other nodes. A temporary panel appears on the bottom, alerting the user that a node has been suggested.

users, in the context of investigative work performed by law enforcement officials, with the aim of facilitating the learning of the use of a new system for the exploration of investigation information.

As explained, the assistive system developed, based on principles from recommender systems and cognitive science, is used in the exploration of a knowledge graph composed of different nodes and connections representing pieces of information collected during an investigation. This knowledge graph implemented here is analogous to one that could be used in a real investigation but customized to provide a goal-oriented task to users: exploring the graph to obtain information necessary to find a target node.

From a technical standpoint, in order to develop such a system, a multitude of components were implemented, as described, including a generator of knowledge graphs, the recommender system itself, a generator of synthetic interaction data used to bootstrap the recommender system, and a generator of recommendations for a given graph based on interaction data (either collected from actual users or generated synthetically).

For the experimental validation of this assistive system, two groups

of participants are proposed: an experimental one, which receives suggestions as needed (based on different criteria), and a control one, which does not receive any assistance from the system. Thus, the expectation would be that participants who receive automated recommendations from the system perform better, in terms of objective metrics (such as performance), implicit features (such as mouse movements), and self-reports (in the questionnaire provided after the main tasks). In addition to an initial validation with civilian participants, a validation should be carried out in collaboration with LEAs and, especially, with the end users of the system.

As mentioned earlier, the system presented in this article works as a web application that runs on a web browser for it to be available online. Due to this, and as an initial implementation of the assistive system, the recommendations were triggered based on different interaction features like the time elapsed, the number of clicks, attempts to solve the task, etc., which were the most viable and appropriate sources, while still providing the necessary information for the knowledge transfer system.

However, more sophisticated methods could be implemented to trigger these recommendations based on the internal states of the users, as inferred in real-time based on various signals. For example, suggestions could be provided based on the estimated cognitive load of the user using pupil dilation signals captured by an eye-tracker, or stress levels could be estimated from physiological signals such as the variation in heart rate using the appropriate sensor.

In conclusion, here, we have proposed a novel assistive system for transferring domain knowledge to novice users, exploiting modern technologies to facilitate the training of users in the use of new digital tools to be used in the field in the course of their work. With this work, the authors would like to highlight how state-of-the-art technologies can be applied by forward-thinking LEAs, with the aim of improving the training and education of current and future law enforcement officials in and for the Digital Age. Furthermore, this work shows of a simple assistive system that learns from previous successful interactions with users could be employed to improve the users' efficacy and efficiency. In Chapter 11, we will explore a more complex system to achieve this, in another data

exploration task using a neuroscientific dataset, and also exploiting physiological signals from the users to achieve a better implicit understanding to deliver timely suggestions.

Chapter 11

EXPERIMENTAL EVALUATION OF AN IMMERSIVE, INTERACTIVE AND CONFLUENT SYSTEM

This chapter is based on:

López-Carral, H., Omedas, P., and Verschure, P. F. (2023c). Experimental evaluation of an immersive, interactive and confluent system. *International Journal of Human-Computer Studies* (submitted)

Generating insights from increasingly large and complex datasets is a challenging task that can be aided by improving the interactions between the users and the systems. Merging advances in psychology's understanding of unconscious processes with the technological advances in virtual reality and sensors engineering, we developed a new 'sentient' system. This 'Sentient Agent' allows for augmented interaction with multiple data sources within an immersive space. By monitoring the users' physiological reactions, it infers their internal states in terms of arousal and cognitive workload and adapts the data presentation according to the perceived

state. Immersion, explicit and implicit interactions were identified as the core features of this experience and were evaluated across three different studies. First, We found that the wider field of view afforded by a large immersive display facilitated recognition and recall of the proposed information and reduced the negative side-effects induced by exploring large data in a limited visual space. Next, explicit interaction was found to be crucial for stimulating a sense of presence and enabled participants to become more involved in their interaction with the content. Finally, the implicit interventions based on the users' unconscious reactions to data helped participants to complete navigational tasks of medium complexity with significantly higher efficiency and efficacy. Furthermore, these interventions provided a more engaging and enjoyable experience. Our work contributes towards the development of a genuinely symbiotic human-computer interaction, in which an immersive and interactive system has the capability of detecting user reactions and using them to enhance the user's experience and performance.

11.1 Introduction

Over the past decades, the amount of digitally collected and stored data has dramatically increased, leading to ever-growing datasets in virtually every field, from science to medicine to financial and business services. In parallel, making proper sense of such large datasets has become a striking challenge. The urge to generate insights and value from these volumes of structured and unstructured data is forcing experts from a multitude of disciplines to rethink their approaches to 'big data'. The conventional method to tackle the current and future 'data deluge' is to analyze and find patterns in complex information using computational techniques from applied mathematics, such as clustering, classification, association rules, and sequential patterns (Tsai et al., 2015). Although these approaches are capable of rapidly processing large amounts of data, addressing the big data problem as only an automated process cannot solve all dimensions of the challenge. Human intervention is still required to obtain meaning-

ful knowledge from data (Mazzocchi, 2015).

A less conventional approach to aid in big-data exploration could be based on exploiting the potential of the user's unconscious processes. The idea behind this would be to bring the human experience to the center of the interaction process. Going beyond traditional interaction, which relies only on the user's explicit input, it is possible to also take into account the implicit states of the users. While explicit inputs correspond to direct actions by the users, such as clicking a button, implicit cues relate to their cognitive or affective conditions. By considering both of these sources of complementary information, a deeper level of integration between humans and computers during the interaction process may be achieved. For example, detecting that the user is overwhelmed by the volume of information presented, a computer system could autonomously reduce the complexity of the proposed content to suit the specific user's needs.

In order to infer internal states of users, a wide variety of techniques have been introduced which rely on the analysis of physiological responses, obtained via different sensors. For example, electrodermal activity and heart-rate variability are well known to correlate with arousal levels (Critchley, 2002; Acharya et al., 2006). Additionally, pupil dilation has been shown to correlate with both arousal and cognitive workload (Bradley et al., 2008; Pomplun and Sunkara, 2003). Recent advances in wearable technologies permit the interpretation of psychophysiological states also in the context of virtual and mixed reality interactions (Betella et al., 2014b).

Here, we introduce an approach based on the combination of virtual reality technologies, multimodal interaction, and implicit feedback to aid in the exploration of large datasets. To validate this approach, we used neuroscientific data as our specific use case. In particular, we developed a tool for the interactive visualization of the human brain Connectome, a large-scale dataset that describes the structural connectivity of an individual brain in the form of a graph.

The final experience, based on the exploration of these neuroscientific data, was deployed in the eXperience Induction Machine (XIM), an immersive space for mixed and virtual reality interaction equipped with sen-

sors and effectors (Bernardet et al., 2010). To assist the user by controlling the interface and guiding the data exploration, we developed an engine integrated into the XIM: the so-called ‘sentient agent’ (SA) (Omedas et al., 2014). The SA is an independent component of the architecture, receiving inputs from the environment and the user (explicit manipulations and implicit reactions) to build a model of the user, which it then uses to adapt the interaction (Wagner et al., 2013).

The resulting system was evaluated to assess the extent to which it supported an enjoyable and effective (efficient, insight generating) exploration of the neuroscience dataset. We identified three core features of the experience, aimed at optimizing it: immersion, explicit interaction, and implicit interaction. Across a series of studies, these independent variables were systematically manipulated to evaluate our approach. Immersion and interactivity were evaluated separately, given that they are two separate variables that must be investigated independently to gain a better understanding of their effect on user experience and performance in virtual environments (McMahan et al., 2006). We hypothesized that explicit and implicit interactions in a highly immersive environment allow for a more engaging, enjoyable, and effective learning experience, than experiencing the same content in a less immersive, passive environment that does not adapt to the user’s subconscious reactions.

In the first study, the effects of being situated in a highly immersive environment (a large projector screen) were compared to those of a less immersive environment (a desktop monitor). In the second study, the effects of interacting explicitly (with gestures and body movements) with a virtual representation of a complex dataset were compared to those of passively watching a pre-recorded, fly-through video of the same dataset. In the final (third) study, the impact on dataset exploration and experience of implicit interaction (i.e., exploration supported and adapted in real-time by the user’s unconscious reactions) was researched. The Goldsmiths’ Research Ethics Committee approved all the studies included in this research.

11.2 Immersion Evaluation

The first dimension of the system that we evaluated was immersion. We sought to explore the impact of deploying our experience in a highly immersive environment versus a less immersive one. To do this, we compared the impact of two vastly different sizes of displays across a series of measures, including recall of the presented information.

11.2.1 Background

Different lines of research suggested that immersive virtual environments can benefit the users, particularly in terms of presence (i.e., the perceptual illusion that a mediated experience is unmediated (Lombard and Ditton, 1997)) and understanding of the explored content.

Compared to smaller screens, large visual displays create a highly immersive experience, eliciting higher levels of realism and superior quality of images, both of which are measures of presence (Baños et al., 2004). Of the many formal features of media, screen size is one of the most systematically studied (e.g., (Lin et al., 2006)). Screen's size influences the way people experience media content, impacting on, for instance, user arousal and enjoyment, and affords a positive media experience overall (Grabe et al., 1999; Lombard and Ditton, 1997). Large projection screens let users build cognitive maps of virtual environments similarly as head-mounted displays (HMDs) (Patrick et al., 2000). High levels of immersion positively affect the task's performance in relation to spatial understanding (Bowman and McMahan, 2007). In their study, that used a visualization system designed for civil engineers, high immersion conditions produced not only faster responses but also responses that were three to ten times more accurate than in the low immersion conditions.

Similarly, immersive virtual environments have been found to provide a higher spatial understanding of complex 3D structures, with experts on underground cave systems answering questions with significantly improved accuracy and speed, and demonstrating greater comprehension compared to a non-immersive environment (Schuchardt and Bowman,

2007). In a different study examining the impact of large high-res displays on data visualization and task navigation performance, it was observed that, with finely detailed data, large high-res displays helped participants find and compare targets faster (up to twice as fast) (Ball and North, 2005). However, higher levels of immersion may not significantly improve performance for activities of a less complex nature, as observed in other studies which failed to find a clear positive relationship between immersion and performance (e.g., (Narayan et al., 2006; Polys et al., 2007)). Additionally, although restricting a viewer's field of view (FOV, which varies according to the actual size of the screen and how far away from the screen the viewer is seated) has been shown to affect memory recall and degrade task performance in real environments (Hagen et al., 1978; Alfano and Michel, 1990), these results have not been replicated in virtual environments (Arthur, 1996).

11.2.2 Objective

The degree of immersion can be objectively assessed as a characteristic of a technology (Slater et al., 1996), and it is defined as the extent to which the display is more or less extensive. The first study was thus designed to test whether the degree of immersion may affect the way people experience visualizations of large datasets. In particular, we hypothesized that the wider field of view afforded by a large display would facilitate recall and recognition of information presented and lead to a higher sense of presence, engagement, and enjoyment, compared to a smaller, less immersive display.

11.2.3 Method

Participants

Forty participants were recruited and paid £10 on completion of the first study. Participants were mainly students from Goldsmiths, University of London. Ages ranged from 19 to 40 years, with a mean age of 24 years ($SD = 4.6$), 50 % females. The majority of participants were British

(62 %), rarely used large screens (such as TVs of 50 or more inches, home projector or cinema) and have no, or very basic, brain anatomy knowledge. None had color vision deficiency or a history of epilepsy. The visuospatial short-term working memory of the participants was tested using an online computerized version of the Corsi Block-Tapping Task (Milner, 1971). In this version of the task, 9 squares are presented in the screen and are illuminated in sequences (of increasing length over the trials) that participants must immediately reproduce by clicking on them in the exact same order (Kessels et al., 2000). There was no significant difference in their Corsi scores between participants allocated to the High and Low Immersion groups, indicating that they were well matched. All participants signed an informed consent form.

Design

Participants were randomly assigned to one of the following two experimental conditions:

- **Low Immersion:** participants experienced a fly-through (pre-recorded video) of a virtual representation of a complex dataset on a desktop display.
- **High Immersion:** participants experienced a fly-through (pre-recorded video) of a virtual representation of a complex dataset on a projection screen.

A passive viewing condition was used in this first study, displaying a pre-recorded interactive exploration of the dataset in order to expose participants to the same content without placing them in the position of interacting with the content in a non-intuitive way.

A between-subjects design was used because participants completing one condition would have been exposed to, and would, therefore, have learned too much about the visualization to be useful in a second condition.

Equipment

For the Low Immersion condition, a 20 inches ProLite E2008HDS Iiyama desktop monitor was used, with a resolution of 1366×768 (see Figure 11.1, left). Conversely, in the High Immersion condition, an Optoma GT1080 projector with a short-throw lens and a light output of 1000 ANSI lumens was used. The video, projected on a 270 cm×220 cm screen (approximately 137 inches), had a resolution of 1920×1080 and a 25,000:1 contrast ratio (see Figure 11.1, right).



Figure 11.1: Setups used in the Immersion Evaluation study. In the Low Immersion condition (left), users were seated in front of a regular computer monitor, while in the High Immersion condition a projector was used for increased screen size.

Procedure

Participants were asked to watch a pre-recorded video of a visual 3D representation of the Connectome dataset for as long as they wished. When participants wanted to stop the video, they could just let the experimenter know. Although participants were instructed to watch the video for as long as they liked (i.e., “until their curiosity was satisfied”), a time limit (of about 12 minutes) was actually in place so that the overall time spent with a participant (for completion of the consent form, watching the video, completing questionnaires, and debriefing) did not exceed 50 min-

utes. This time limit enabled researchers to allocate one-hour time slots for all participants.

The 12-minute video produced for this experiment provided participants with a 360-degree panoramic view of the 3D brain visual representation of the Connectome dataset, representing the structural connectivity of an individual brain in the form of a graph, rotating in a loop. The video showed the brain model slowly rotating in a loop, revealing five areas (i.e., Frontal, Parietal, Temporal, Occipital, and Cingulate). Each of these areas contained a large number of Region of Interests (or ROIs, i.e., units of neuronal groupings), represented as nodes of different colors, and interconnections between nodes, represented as edges (see Figure 11.2).




Figure 11.2: 3D visual representation of the Connectome dataset used in the current study. This large-scale dataset describes the structural connectivity of an individual brain in the form of a graph. Different regions of interest (brain areas) are represented as colored nodes and their interconnections as lines. Participants were asked to watch a video showcasing the dataset for up to 12 minutes.

A consistent experimental procedure was used in both conditions. Participants were seated at a standard distance appropriate to the screen size provided. For the Low Immersion condition, the monitor was positioned at a comfortable distance of 90 cm from the participant, giving a viewing angle of 27.6 degrees. For the High Immersion condition, the distance to the screen was 230 cm, giving a viewing angle of 66.8 degrees.

Measures

After the experimental session, participants filled in a short version of the ITC-Sense of Presence Inventory (SOPI) (Lessiter et al., 2001), which produces four scores related to a media experience: sense of physical space (i.e., presence or “being there”); engagement (i.e., psychologically involved); ecological validity (i.e., lifelike or real); and negative effects (i.e., adverse physiological reactions such as dizziness or nausea). The 10 items were rated on a 1-5 scale from “strongly disagree” to “strongly agree”. Additionally, two sub-scales of the Intrinsic Motivation Inventory (IMI) (Ryan, 1982) were used to measure interest/enjoyment and value/usefulness (14 items). Response options ranged from 1 “not at all true” to 7 “very true”.

Finally, memory recall and recognition tasks were performed by participants, who were asked to recall different node colors and assign them to the correct areas; recognize areas with the fewest and greatest number of nodes, and which among three pictures best matched the environment previously displayed. Additionally, to categorize whether memory was intentional and incidental, the following question was also posed: “during the exploration, did you focus on any of the information below with the intention of memorizing it? [colors of nodes, association of these colors of nodes with different areas, number of nodes in relation to different areas, the shape of the Connectome]”. The memory questionnaire was scored by assigning one point to each correct answer. As such, the scores ranged from 0 to 15.

Time spent by participants watching the video was recorded during the experimental session and used as an objective measure of enjoyment and engagement (Guo et al., 2014; Park et al., 2016; Wu et al., 2018).

11.2.4 Results

A series of independent samples *t*-tests were performed to test for differences between the two groups. Participants in the High Immersion condition performed significantly better in the recall and recognition tasks

($M = 10.05, SD = 2.98$) than participants in the Low Immersion condition ($M = 8.20, SD = 2.46$); $t(38) = 2.14, p = 0.039$ (see Figure 11.3(a)). When explored in more detail, analyses on the recall and recognition results revealed that participants in the High Immersion condition reported higher scores on all memory items. In particular, the following revealed significant differences between the conditions:

- Recalling colors of nodes: participants in the High Immersion condition ($M = 6.10, SD = 2.40$) performed significantly better than participants in the Low Immersion condition ($M = 4.45, SD = 2.04$); $t(38) = 2.34, p = 0.025$.
- Matching the color of the nodes to the correct areas: participants in the High Immersion condition ($M = 2.1, SD = 1.83$) performed significantly better than participants in the Low Immersion condition ($M = 1.0, SD = 1.16$); $t(38) = 2.263, p = 0.031$.
- Recognizing the shape of the Connectome: a chi-square test of independence revealed that participants in the High Immersion condition performed significantly better than participants in the Low Immersion condition; $\chi^2(df = 1) = 4.402, p = 0.036$.

An independent samples t -test was also performed to detect mean differences between groups on successfully recalled features by incidental (i.e., recall of information not specifically attended to) and intentional (i.e., recall of information which participants focused on) memory. However, no significant differences between the two groups were found on scores for intentional or incidental memory. Participants in the Low Immersion condition ($M = 1.35, SD = 0.58$) also reported a significantly higher level of negative effects than participants in the High Immersion condition ($M = 2.05, SD = 1.36$); $t(38) = -2.19, p = 0.028$ (see Figure 11.3(b)). No significant differences between the High and Low Immersion conditions were found for spatial presence, engagement, ecological validity, enjoyment/interest, value/usefulness, or time spent watching the video.

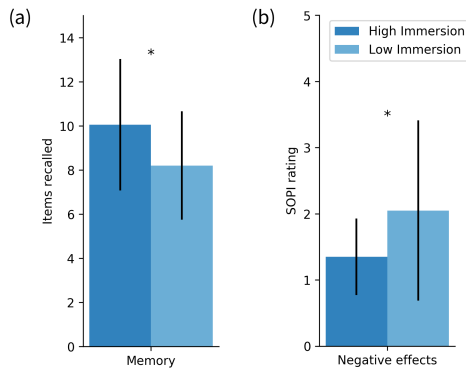


Figure 11.3: Differences between the High and Low Immersion conditions in the Immersion Evaluation. **(a)** Participants in the High Immersion condition recalled a significantly higher amount of information than those in the Low Immersion condition. **(b)** Participants in the High Immersion condition reported a significantly lower amount in the negative effects part of the SOPI than those in the Low Immersion condition.

11.2.5 Discussion

In the Immersion Evaluation, we wanted to study the effect of two different degrees of immersion during the visualization of large datasets, in terms of subjective experience and value obtained from the data. We hypothesized that with higher levels of immersion, afforded by a large display covering wide field of view, participants would have a more fruitful experience, as measured by a series of tests and questionnaires. Indeed, we found that participants who watched the video on a large immersive screen recalled and recognized more features of the Connectome dataset than those who viewed the data on a small screen. Although the sense of presence and engagement reported by participants in the High Immersion condition were not significantly higher than those of participants in the Low Immersion condition, the difference in screen size and viewing angle affected the way participants in the two groups experienced the visualization of a large dataset. The results support the hypothesis that immersion affects the amount of information stored in the individual's

memory. These results are in line with a previous study that reported that subjects were most effective at forming spatial knowledge of abstract information when it was presented on a large display in comparison to a small one (although this effect was also dependent on the higher resolution of the display) (Ni et al., 2006). Likewise, another study suggested that an increase in participants' ability to memorize relates to the richness and quality of the spatial cues provided by the virtual environment (Ragan et al., 2010). The latter was also confirmed in a separate study, showing that the more detailed the task, the higher the level of immersion needed in order to learn and recall information effectively (Sowndararajan et al., 2008). The authors concluded that "in the high-immersion condition, the objects in the environment were spread out spatially [...] and could be remembered based on their spatial location, while in the low-immersion condition, all of the objects were 'squeezed' into a much smaller physical space". Additionally, immersion was found to be negatively correlated to negative effects. Watching the video on a large screen was found to reduce side-effects, such as dizziness, disorientation, or tiredness, most probably induced by exploring large data in a limited visual space.

11.3 Explicit Interaction Evaluation

The second dimension of the system that we evaluated was explicit interaction. We sought to explore the impact of allowing participants to actively interact with the dataset that was presented versus passively observing it. Based on the results of the previous evaluation, this was performed using a large projection screen.

11.3.1 Background

A core feature of the developed system considers the delivery of 'presence', not only in terms of letting users to 'step inside' large datasets but also by enabling them to interact with the content intuitively. It has been shown that, when users test body movements and actions in an inter-

active environment (e.g., virtual world) and the system reacts following the users' expectations (i.e., affordance), the trust in that environment is reinforced (Lombard and Ditton, 1997). This suggests that, if interactive controls are intuitive enough, users may forget about the 'controller', supporting a transparent way of interaction with the virtual world. Under these conditions, it is predicted that people are less likely to perceive the mediated experience as unmediated, i.e., sense of presence (e.g., (Bayliss, 2007; Crick, 2011)). Input devices that enable interaction through gestures and body movements (such as keyboards, joysticks, 3D mice and motion sensing input devices), give the user 'perceived control' of the displayed environment which may not only enhance the sense of presence (Hendrix and Barfield, 1996), but may also provide for a more enjoyable experience (e.g., (Shim et al., 2003)). Collated evidence also suggests that students learn more effectively when they are able to interact with learning materials, as the manipulation of data encourages a more active search for meaning than with direct instruction (Moreno, 2005). The enjoyment experienced as a result of an interactive exhibit in the Science Museum of London was, for example, found to translate a learning experience into a form of play (Haywood and Cairns, 2006).

11.3.2 Objective

In our second study, we aimed to test whether the gesture and body-based user interaction can affect the way people experience the visualization of large datasets. In particular, it was hypothesized that interacting with the Connectome dataset would facilitate better recall and recognition of the information presented and lead to a higher sense of presence, engagement, and enjoyment, compared to passively watching a pre-recorded, fly-through video of the same dataset.

11.3.3 Method

Participants

As with the first study, forty participants were recruited and paid £10 on completion of the study. Participants were mainly students from Goldsmiths, University of London. Ages ranged from 19 to 46 years, with a mean age of 24 years ($SD = 7.1$), 50 % females. The majority of participants were British (55 %), rarely used large screens and have no, or very basic, brain anatomy knowledge. None of them had color vision deficiency or a history of epilepsy. The groups did not differ significantly (demographically and with regards to their working memory), indicating that they were well matched. All participants signed an informed consent form.

Design

Participants were randomly assigned to one of the following two experimental conditions:

- **Explicit Interaction:** participants actively interacted using gestures and body movements with the Connectome dataset displayed on a large projection screen;
- **No Interaction (or yoked condition):** each participant in the no interaction group watched a fly-through video from a matched paired participant in the Explicit Interaction condition. The fly-through of the Connectome dataset was presented on the large projection screen. For example, participant A in the No Interaction condition watched the recorded video of participant A in the Explicit Interaction condition; participant B in the No Interaction condition watched a pre-recorded video of participant B in the Explicit Interaction condition; and so on.

As in the Immersion Evaluation, a between-subjects design was used because participants' learning outcomes in the second condition of a repeated measures design would be contaminated by their experience of the

content in the first condition. Pairing participants in the No Interaction condition to those in the Explicit Interaction condition was chosen over a pre-recorded video used in the first experiment in order to expose participants to exactly the same content.

Equipment

For the Explicit Interaction condition, a Microsoft Kinect 2.0 was used to recognize participants' hand gestures and body movements, which allowed them to zoom in and out of the Connectome 3D model, rotate it up and down, and find more information about its different areas. A screen capture tool (Bandicam, recording at 60 fps, resolution 1920×1080, XVID codec, 2.3 Mbps) was used to record the interactive experience. The same projector used in the Immersion Evaluation (270 cm×220 cm screen, 1920×1080 resolution, 25,000:1 contrast ratio) was used to visualize the dataset in both conditions.

Procedure

A consistent procedure was used for both experimental conditions. Participants were asked to stand at 230 cm from the screen, as per the High Immersion condition in the Immersion Evaluation study. This is the distance at which the Kinect does not recognize the user in front of the sensor to be moving. Participants were then taken through a short and straightforward practice training session, which helped them become familiar with the gestural control and reduce the novelty effect on performance. Participants were asked to interact with a version of the Connectome environment in which nodes and edges were presented in a random order (i.e., neutral content). The experimenters ensured that all participants were able to complete each action before proceeding to either the 'interaction' or 'fly-through video' phase that followed. The training session was performed by both groups to control for the effect of interacting with neutral content before the experiment. Training time was recorded to control for the possibility of differences in interaction time affecting the results.

The Connectome dataset explored in the experimental session was the same used in the Immersion Evaluation. However, rather than watching a pre-recorded video of the dataset, participants in the Explicit Interaction condition were asked to interact with the virtual brain and explore all of its areas for as much time as they like (but limited to about 20 minutes). When participants wanted to stop the video, they could just let the experimenter know. Each participant in the No Interaction condition was paired to and watched the recorded video of a participant in the Explicit Interaction condition. As with the Explicit interaction condition, no time limit was imposed, but the total duration of the video was dependent on the duration of the interaction of paired participants.

Measures

The same measures used in the Immersion Evaluation were used for this study: a short version of the ITC-SOPI to rate presence, engagement, ecological validity, and negative effects; two subscales of the IMI to measure interest/enjoyment and value/usefulness; and a recall test to compare the memory performance between conditions, together with the time spent.

11.3.4 Results

An independent samples *t*-test was conducted to test whether there were any significant differences in mean scores for spatial presence between the two groups. Results revealed that participants in the Explicit Interaction condition reported a significantly higher sense of presence ($M = 3.62, SD = 0.67$) than participants in the No Interaction condition ($M = 3.00, SD = 0.76$); $t(38) = 0.66, p = 0.010$ (see Figure 11.4(a)). A Spearman's rank-order correlation further revealed a strong, significant positive correlation between presence and enjoyment $r(80) = 0.513, p = 0.0008$. A significant difference in enjoyment/interest scores was also found, $t(38) = 0.25, p = 0.020$ with participants in the Explicit Interaction condition ($M = 5.07, SD = 1.28$) reporting higher scores than participants in the No Interaction condition ($M = 3.97, SD = 1.57$)

(see Figure 11.4(b)). This was reflected in the time spent (in minutes) by the two groups in exploring the dataset, with participants in the Explicit Interaction engaging for a significantly longer time ($M = 5.91, SD = 1.80$) than those passively watching the video ($M = 3.58, SD = 1.83$); $t(23) = 3.06, p = 0.005$ (see Figure 11.4(c)). Additionally, participants in the Explicit Interaction condition reported a significantly lower level of negative effects ($M = 1.60, SD = 0.94$) than did participants in the No Interaction condition ($M = 2.40, SD = 1.39$); $t(38) = -2.13, p = 0.040$ (see Figure 11.4(a)). No significant differences between the Explicit and No Interaction conditions were found for performance (measured by the number of items recalled and recognized), engagement, ecological validity, or value/usefulness.

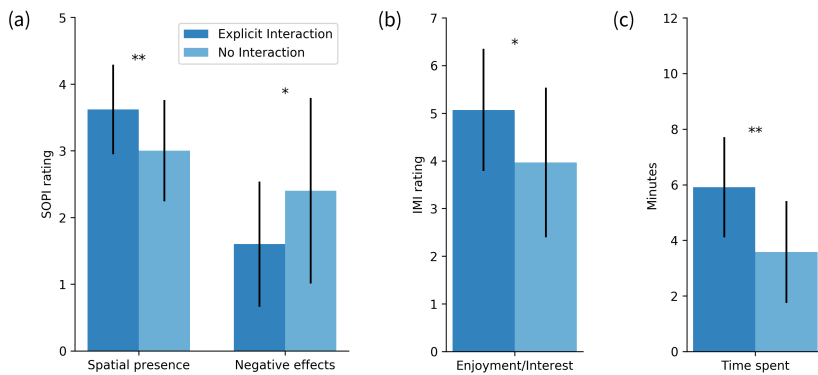


Figure 11.4: Differences between the Explicit Interaction and No Interaction conditions in the Explicit Interaction Evaluation. **(a)** Participants rated a higher spatial presence and fewer negative effects when they were able to interact with the content in comparison to when they could not. **(b)** Participants reported a significantly higher level of Enjoyment/Interest in the experience in the Explicit Interaction condition. **(c)** Participants spent a significantly higher amount of time exploring the dataset when they were able to interact with it.

11.3.5 Discussion

In the Explicit Interaction Evaluation, we wanted to assess the effects of interacting with large datasets using gestures and body movements. We hypothesized that directly interacting with the visualization would result in a better experience that leads to better knowledge extraction, compared with passively watching a pre-recorded video of the same dataset. Results from this study indicated that participants who interacted with the Connectome dataset using a large immersive screen did not recall nor recognize more features of the dataset than those who just watched a fly-through video. However, explicit interaction elicited a significantly higher sense of presence and enjoyment, as well as fewer adverse effects. The results also suggested that as the sense of presence increased, the level of enjoyment experienced by users increased. Enjoyment was reflected in the significantly longer time spent by the Explicit Interaction group in exploring the dataset. Compared to passive viewing in an immersive environment, gesture-based environment interactions increased both the sense of ‘being there’ and positive experiences of that environment. Similar results were also observed in a study on virtual museum experiences (Sylaiou et al., 2010). Further, the negative effects reported by participants in the No Interaction condition support other findings in the literature. While investigating motion sickness in virtual environments, it was found that a lack of control over their movements induced a higher level of sickness symptoms in the subjects. These results have been explained in the context of the sensory conflict theory (Reason and Brand, 1975), where the participant is more likely to experience sickness due to an unexpected conflict occurring between sensory inputs (passive viewing) (Sharples et al., 2008).

11.4 Implicit Interaction Evaluation

The final dimension of the system that we evaluated was implicit interaction. We sought to explore the impact of receiving assistance from the system itself based on the internal states of the users, as inferred dynam-

ically. Expanding on the results of the two previous evaluations, participants navigated the system in an immersive environment and actively interacting with the information presented.

11.4.1 Background

Implicit interaction, which takes into account the user's unconscious reactions in addition to explicit input, encourages a more symbiotic relationship between humans and computing devices, a synergy that has recently been labelled with the term Human-Computer Confluence (HCC). As an emerging field of research, HCC applies to the new classes of user interfaces that "make use of several sensors and are able to adapt their physical properties to the current situational context of users" (Ferscha, 2013). Within these confluent systems, implicit interaction has become the invisible interface through which a user can unconsciously influence the behaviour of the system and what is being presented.

The importance of integrating user psycho-physiological responses and emotions into the machine environment was first emphasized by Picard in 1995, who proposed that computers that interact naturally and intelligently with humans need the ability to recognize and express affection. The author's idea was that it should be possible to create machines that relate to, arise from, or deliberately influence emotion or other affective phenomena (Picard, 1995). The author coined the term "affective computing" to define her idea, which provides a glimpse of how vast this domain is. The variety of possible roles and functions that affective considerations introduce into the relationship between humans and machines is extensive. It can range from "recognizing user affect" and "adapting to the user's affective state", to "generating 'affective' behaviour by the machine, modelling user's affective states, or generating affective states within an agent's cognitive architecture" (Hudlicka, 2003). Although initial efforts were predominantly focused on emotion detection and recognition, recent studies on affective computing have begun to explore how to exploit emotional information to enhance the user experience by adapting software in real-time (Aranha et al., 2019). This approach has been

applied to many different scenarios: e-learning systems, in which students may receive personalized support by detecting the learners' affective states (Santos, 2016); video games, in which different game's features may adapt based on the emotions of the player (Bontchev, 2016); social robots to assist people with special needs in more humane ways (Magnenat-Thalmann and Zhang, 2014).

Research in HCC has also been focused on brain-computer interfaces (BCIs) (Gaggioli et al., 2016), to allow people with severe motor disabilities to interact with computers through their brain's activity without requiring physical action (e.g., selecting a letter from a virtual keyboard or moving a robotic device). This kind of brain-computer interfaces, consciously controlled by the user, has been defined as 'active BCI' (Zander et al., 2010), to distinguish them from those that consider brain activity as an additional source of information used to augment and adapt the interface (Girouard et al., 2013). The author referred to this latter type of BCI as 'passive'. Passive BCI targets a broader group of users for whom current active BCIs are impractical because of the amount of attention and concentration required to the user to interact with the machine actively (Zander et al., 2010). Passive BCI has been successfully applied to recognize dominant emotions and automatically classify, or 'tag', visual multimedia content (Yazdani et al., 2009). More recently, studies on driving assistance applications have also explored the use of passive BCIs to predict the driver's steering intentions and subsequently trigger driver support systems for increasing traffic safety and avoiding fatalities. Similarly, a system has been proposed that identifies high mental workload in drivers operating under real traffic conditions (Kohlmorgen et al., 2007). This information is then used in real-time to mitigate the high mental workload induced by the influx of information generated by the car's electronic systems. A different area of research explored the use of subliminal cueing in virtual environments revealing that non-invasive wearable sensors designed to support users in complex 'real world' activities through subliminal visual and audio cues are highly effective, with subliminal cueing significantly improving performance, and overall user experience (DeVaul et al., 2003).

Research on affective computing, BCI technology, and, more recently, on HCC have contributed to elucidate the emerging symbiotic relationship between humans and machines and to the development of novel adaptive systems that can perceive, react to, and even affect our mental state in a meaningful way. Integrating information on cognitive aspects of user state into computers may lead to more natural interactions with the real or virtual world.

11.4.2 Objective

In this last, third study, we investigated the benefits of implicit interaction. Experimental conditions with and without implicit interaction were evaluated and compared using both objective (e.g., accuracy, speed) and subjective measures. In particular, it was hypothesized that our Sentient Agent, which monitors the user's mental state and optimizes the presentation of data to the user's state (e.g., by increasing or decreasing the salience and the number of visual information presented to the users depending on their cognitive workload and arousal), would facilitate the user's navigation performance and interaction in the virtual environment.

11.4.3 Method

The experimental methodology followed for this study is also reported in (Cetnarski et al., 2015).

Participants

Fifty-one participants were recruited and paid 10€ on completion of the study. The study was conducted at the Universitat Pompeu Fabra in Barcelona (Spain), which hosts the eXperience Induction Machine (XIM), the mixed reality environment used in this project and equipped with sensors monitoring the user's body movements and psycho-physiological signals (Bernardet et al., 2010). Ages ranged from 18 to 39 years, with a mean age of 22 years ($SD = 4.3$), 45.54 % females. The majority of

participants were students from UPF, Spanish (78 %), rarely used large screens and had no, or basic brain anatomy knowledge. None of them had color vision deficiency or a history of epilepsy. As in the Immersion Evaluation and the Explicit Interaction Evaluation, the groups did not differ significantly on their visual short term working memory tested using the Corsi Block-Tapping Task. All participants signed an informed consent form. All participant forms (consent form, instructions, questionnaires, and debriefing form) were translated into Spanish by a certified translator, so that native Spanish speakers could be recruited for this study.

Design

Participants were randomly assigned to one of the following three experimental conditions:

- **Congruent Implicit Interaction:** in this condition, the SA monitored both explicit and implicit signals acquired from the users and optimized the data presentation according to feedback from both of these types of reactions when most needed by users;
- **Incongruent Implicit Interaction:** the SA monitored both explicit and implicit signals received from the users and optimized the data presentation according to feedback from both of these types of reactions when least needed by users;
- **No Implicit Interaction:** this was the Control condition in which the SA monitored both explicit and implicit signals received from the users, but only the explicit signals (i.e., hand gestures, body movements) were used to control interactions with the content presented.

‘SA turned on’ is the term hereby used to describe the situation in which the SA was actively intervening when some criteria relevant to the implicit state of the participants were met in the Congruent and Incongruent condition. The Incongruent condition was added to test whether

the SA settings selected for interventions were correct (see section 11.4.3 for more details). This was different from the Control condition (i.e., No Implicit Interaction), in which the SA was turned off, meaning that the SA was not actively intervening, although the user states were still being monitored.

As in the Immersion Evaluation and the Explicit Interaction Evaluation, a between-subjects design was used. A within-subjects variable was added to this design, which consisted of three levels (i.e., easy, medium, and hard) of graph complexity. This variable was added to test how different degrees of complexity of the network of nodes and lines presented to the user may affect performance.

Equipment

A range of custom-made wearable sensors was used to monitor participants' (implicit) physiological responses (Wagner et al., 2013): a sensing glove that monitored skin conductance (Tognetti et al., 2007; Carbonaro et al., 2012); a chest band that monitored heart rate and respiration (Paradiso et al., 2005); and an eye tracker which was used to monitor changes in pupil size (Lanatà et al., 2011). Sensors were worn by participants in all conditions to avoid any possible influence of such components on the experiment results. Participants' hand gestures and body movements were captured by Microsoft Kinect 2.0.

Procedure

Participants in all conditions were initially asked to stand straight and still on a fixed position in the middle of the room (marked with an 'X' on the floor) whilst the experimenter checked that the sensors were working. For the remainder of the study, they were free to move closer to and farther away from the screen to zoom in and out, respectively (see Figure 11.5). Similarly to the Explicit Interaction Evaluation, participants were taken through a short and simple practice training session, which helped familiarize them with the gestural control.




Figure 11.5: Conceptual representation of a user interacting in the XIM. Participants are free to move around the space to explore the dataset in an embodied manner.

Free exploration of the dataset was considered inadequate for the needs and goal of the Implicit Interaction Evaluation. In order to systematically test whether implicit interaction was beneficial, a specific goal was introduced, so that the potential added value of implicit interaction could be measured in terms of how adequately the SA supported users in achieving their goal. Therefore, an ‘artificial’ dataset was chosen over the Connectome dataset used in the first two experiments. The artificial dataset consisted of a network of nodes interconnected by lines. Two of these nodes were of different colors to represent the starting (yellow) and ending (blue) points. Participants’ objective was to reach the ending node (or target) following the shortest path. At every step, participants were presented with three (white) nodes to choose from and had to select the one that would bring them closer to the target. They had up to 30 seconds to make a decision and select a node. A timer appeared in the upper left-hand corner to remind them how much time they had left. A score was provided at the end of each trial to gamify the experience and keep the participants engaged: participants got 10 points for each correct node and lost 20 points for each wrong node. Figure 11.6 shows an example of a path followed by a participant during a task.

On arrival at the target node, participants were presented with the same pattern of nodes again, with the same starting and ending point.

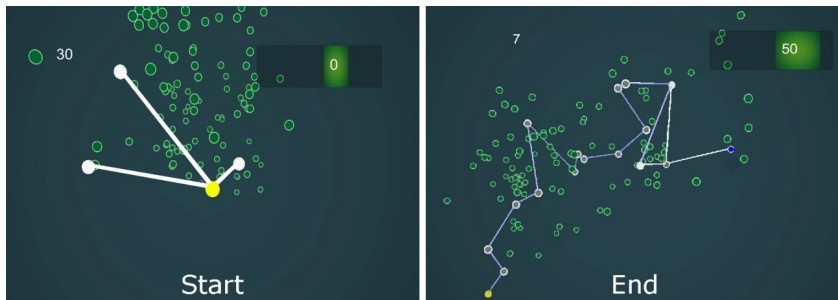


Figure 11.6: Screenshots of starting (yellow) and ending (blue) points in the ‘hard’ level of graph complexity at the start (left) and at the end (right) of the task. The number on the top left of both screens is the timer. The number on the top right of both screens (on the green background) is the score meter.

Learning from their experience with the first trial, they had a second opportunity to reach the target in the fewest steps possible. After completion of this second trial, a new network of nodes of different difficulty was presented. Participants were instructed to look for the target node and reach it, again in the fewest steps possible and, once they completed this, they then had a second attempt. Finally, the network was changed for the last time, and participants had to repeat the task for the last two times. Participants in the Congruent and Incongruent conditions were instructed that the system may intervene to aid them in their decision. When this happens, they will see a dark blue horizontal stripe on the top of the screen. This instruction was added to make them aware when possible sudden changes occurred and thus prevent them from stopping the exploration.

The three different networks were presented (i.e., easy, medium, hard) in a randomized order, and each was presented twice in succession, giving a total of 6 trials per participant. Each network was composed of a fixed number of steps ranging from 12 (i.e., the fewest number of possible nodes selected) to 36 (i.e., the highest number of possible nodes selected). For the Congruent and Incongruent condition, the SA was ‘turned on’ in the first trial of each network, and turned off in the second trial. The

Table 11.1: Experimental design of the Implicit Interaction Evaluation

	Network difficulty level: Easy		Network difficulty level: Medium		Network difficulty level: Hard	
	1 st trial	2 nd trial	1 st trial	2 nd trial	1 st trial	2 nd trial
Control Condition	SA	SA	SA	SA	SA	SA
SA turned off	off	off	off	off	off	off
Congruent Condition	SA	SA	SA	SA	SA	SA
SA turned on when most needed by users	on	off	on	off	on	off
Incongruent Condition	SA	SA	SA	SA	SA	SA
SA turned on when least needed by users	on	off	on	off	on	off

second trial was used as a direct measure of how SA interventions in the first trials helped participants performance in the second trial. Participants wore sensors in all conditions to avoid any bias in the experiment results. The experimental design for this study is summarized in Table 11.1.

Criteria for Sentient Agent Interventions

In order for the SA to intervene in the Congruent and Incongruent conditions, a set of criteria was defined. The first criterion was based on the user’s accuracy in node selection: the SA intervened only when a participant selected a wrong node. This criterion was adopted to avoid countless SA interventions from occurring unnecessarily. This would have created a rather confusing experience, as sometimes participants were performing correctly regardless of SA interventions. SA’s goal was instead to convey a seamless experience in which participants only received help

when needed. The intervention was set to last until participants made two correct choices in a row after a wrong selection.

The second criterion concerned the user's state analyzed by the SA once a wrong node was selected. This differed for the Congruent and Incongruent conditions. For the Congruent condition, the SA intervened when arousal and/or cognitive workload increased relative to baseline. Conversely, participants in the Incongruent condition received the SA support when arousal and/or cognitive workload decreased.

The cognitive workload was estimated from the pupil dilation of the user, obtained from the wearable eye tracker (Betella et al., 2014a). A baseline of pupil size was first taken within a time window of five seconds. A shorter stimuli-related window of two seconds was then compared to this baseline creating a moving average of pupil size based on both the long and short time windows. In this way, the system identifies the increases in pupil size resulting from sustained task-related cognitive activity enabling a calculation of the 'cognitive load'. The moving average method was also used to estimate user's arousal from the electrodermal activity (EDA) and the heart-rate variability (HRV). A baseline was established through a longer 30 second time window while a shorter window of 10 seconds was used to gather a measure of arousal. The calculated changes in the signal feedback from EDA and HRV were combined to create a weighted average representing the user's state of arousal.

Types of Sentient Agent Interventions

SA's interventions were based on the user's physiological signals indicating their underlying internal states related to cognitive workload and arousal.

Cognitive workload theory proposes attention as a selective process that determines what information gains access to conscious awareness (Lavie et al., 2004). However, during this selection process, most of the information is filtered out from subjective experience and replaced by internally generated estimates or expectations. Under increased cognitive workload, such filtering and substitution processes become more

frequent, which can lead to the omission of relevant information (Mathews et al., 2014). A possible countermeasure for such stimuli omission is to increase the salience of relevant information, thus creating a discrepancy between perceived and expected stimuli. For this reason, cognitive workload was associated with the level of visual information that participants were able to process and understand. When a participant's cognitive workload increased (for Congruent) or decreased (for Incongruent), the SA intervened by:

1. decreasing the complexity of the network (e.g., the number of nodes was reduced to make the network clearer and less complex), and
2. reducing the salience of one of the two incorrect nodes (e.g., making the node more transparent and not selectable).

Previous studies have demonstrated that visual motion of a stimulus affects the user's internal state, and, in particular, arousal (Detenber et al., 1998; Simons et al., 2000). In order to maintain a balance in the user's arousal levels, one of the possible SA interventions consisted of modulating the interaction speed. Low arousal levels were counterbalanced with an increase in speed, while in the presence of high arousal, the SA slowed down the pace of the interaction. Additionally, acoustic features such as pitch can convey emotions (Burkhardt and Sendlmeier, 2000). Previous research has indeed demonstrated that pitch can be effectively used in the sonification of complex data in a wide range of tasks (Walker, 2002; Flowers, 2005). For this reason, as an additional intervention to modulate arousal we employed auditory cues corresponding with different levels of pitch.

Arousal was used by the SA to assist the user interaction with the dataset. When arousal increased (for Congruent) or decreased (for Incongruent), the SA intervened by:

1. reducing the speed of the interaction (e.g., making the pointer less sensitive to user gestures and the embodied navigation slower), and

2. adjusting the sonification parameters to induce changes in the arousal and offering a guiding panning cue (e.g., the volume was increased on either the left or right side depending on where the correct node was located).

These interventions could occur independently of one another: sometimes participants only received interventions triggered by an increase or decrease in arousal in either the Congruent or Incongruent conditions, respectively. At other times participants only received interventions triggered by an increase or decrease in cognitive workload in the Congruent or Incongruent conditions, respectively. When arousal and cognitive workload were both either increasing or decreasing, all interventions occurred simultaneously. See Figure 11.7 for an overview.

Measures

As with the Immersion Evaluation and the Explicit Interaction Evaluation, a short version of the ITC-Sense of Presence Inventory and two subscales of the Intrinsic Motivation Inventory measuring interest/enjoyment and value/usefulness were used.

The effectiveness of implicit interaction was evaluated using a series of objective measures, namely: a) the number of steps taken to reach the target node, b) the time taken to reach the target node, and c) the accuracy in reaching the target node. Accuracy was calculated as the proportion of times participants in the Control condition made the right choice divided by the total number of nodes selected by these participants. The number of times a correct selection occurred during an SA intervention was then assessed for the Congruent and Incongruent conditions and divided by the total number of nodes selected by each of these two groups during SA interventions. The proportion of times participants made the right choice in the Control condition was thus compared to the proportion of times the SA helped participants made the right choice in the Congruent and Incongruent conditions.

Finally, to test whether SA Congruent interventions during the first trial helped participants to perform better in the second trial, only those

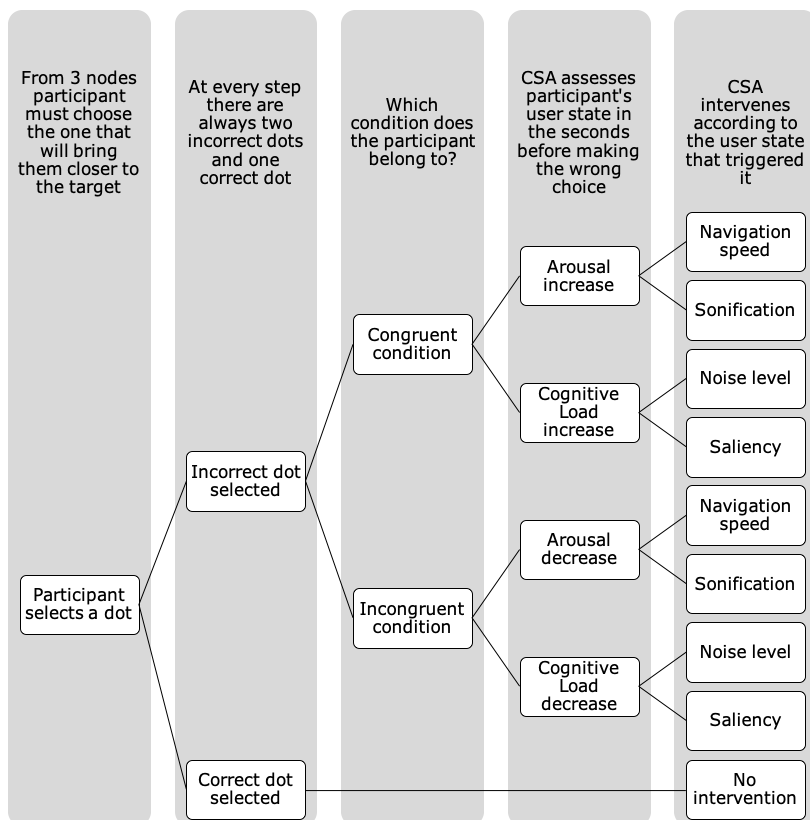


Figure 11.7: Overview of the types of SA interventions triggered by different user states: interventions such as changes in network complexity and node salience occurred with an increase or decrease of cognitive workload, while changes in interaction speed and sonification occurred with an increase or decrease in arousal.

nodes that were selected in both first and second tasks by participants were analyzed. This was obtained by dividing the total number of times a correct selection occurred in the first trial by the total number of SA interventions that occurred in the first trial. The result was labelled ‘memory accuracy’, i.e., the proportion of times that SA interventions in the first trials helped participants perform better without the SA in the second tri-

als.

11.4.4 Results

Frequency of Sentient Agent Interventions

For one of the participants, the physiological responses were not recorded due to a software error, and they were therefore excluded from the analysis. The remaining 50 participants considered in the analysis completed six trials each (an initial trial and the second trial of each of the three levels of difficulty), giving a total of 300 trials completed overall. Overall, the number of SA interventions was equally split between the Congruent and Incongruent conditions (Figure 11.8), with a higher number of interventions for the medium and the hard difficulties. SA interventions only occurred in the first trial of each level of difficulty. For this reason, when comparing the number of SA interventions between the Congruent and Incongruent conditions, only the first trials were considered in the analysis.

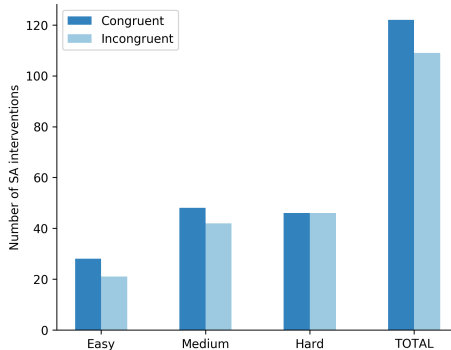


Figure 11.8: Visualization of the number of times the SA was triggered by an increase (Congruent) or decrease (Incongruent) in arousal and cognitive load (both variables appear together here) by level of difficulty (easy, medium and hard).

SA interventions occurred at least once on 57 of the 102 total first

trials (i.e., 56 %). The SA intervened 231 times overall, with an average of 2.5 interventions per task, and this frequency was distributed among the three levels of difficulty. The SA intervened 49 times overall in the easy level of difficulty, 90 times in the medium and 92 times in the hard.

Objective Measures of Performance

In order to investigate the effects of the SA on participants' performance, only the first trials in which participants in the Congruent or Incongruent condition reached the target by selecting at least one incorrect node were considered (which meant that the SA intervened at least once). For consistency, any first trial completed without mistakes by participants in the Control condition was also excluded. As a result, of the 150 total first trials, 55 were excluded from the analysis.

A series of one-way ANOVAs were performed to examine, for each individual level of difficulty, potential differences across the three conditions on different objective measures of performance. Outliers were defined as values over 1.5 times the interquartile range away from the median. Identified outliers were removed from the data before statistical analysis of the differences between groups.

ANOVAs were performed to determine whether mean scores of the number of steps taken (ranging from 12 to 16 nodes) to reach the target node in the first trials of each level of difficulty were significantly different across the three conditions (i.e., Control, Congruent, Incongruent). Results from these analyses revealed that there were no significant differences between groups for the easy and hard levels of difficulty. However, statistically significant differences were found for the medium level of difficulty; $F(2, 28) = 12.78, p = 0.00$. A Bonferroni post-hoc test revealed that participants in the Congruent condition ($M = 12.5, SD = 0.70$) performed significantly better than participants in both Control ($M = 14.7, SD = 0.71$) and Incongruent ($M = 13.7, SD = 1.22$) condition. This suggests that Congruent SA interventions account for a 15 % reduction in the number of steps taken to reach the target node (i.e., a performance benefit) in the medium level of difficulty compared to the Control

condition, and for a 9 % compared to the Incongruent (see Figure 11.9(a)).

A statistically significant difference was also found in the time (in seconds) taken to reach the target node in the first trials in the medium level of difficulty [$F(2, 33) = 5.69, p = 0.008$]. Because participants could spend as much time as they wanted on the first node, this time was excluded by the analysis. A Bonferroni post-hoc test revealed that participants in the Congruent condition ($M = 99, SD = 25.72$) performed significantly better than participants in the Incongruent condition ($M = 166, SD = 77.96$) with a reduction in time taken of approximately 40 %, but not better than the Control condition ($M = 112, SD = 29.98$) (see Figure 11.9(b)). Additionally, statistically significant differences were found between Incongruent ($M = 101, SD = 49.97$) and Control ($M = 164, SD = 42.67$) conditions in the hard level of difficulty [$F(2, 31) = 4.06, p = 0.027$], but not between these conditions and the Congruent condition ($M = 192, SD = 108.12$).

A series of one-way ANOVAs were also performed to determine whether total scores of nodes correctly selected in the first trials were significantly different across the three conditions (i.e., Control, Congruent, Incongruent). Results from these analyses revealed that there were no significant differences between groups for easy and hard level of difficulty. However, statistically significant differences were found for the medium level of difficulty [$F(2, 27) = 4.07, p = 0.028$]. A Fisher's Least Significant Difference (LSD) post-hoc test revealed that participants in the Congruent condition ($M = 0.97, SD = 0.83$) performed significantly better than both participants in the Control ($M = 0.80, SD = 0.57$) and Incongruent condition ($M = 0.84, SD = 0.19$) (see Figure 11.9(c)).

Another hypothesis that was tested was whether SA Congruent interventions in the first trial helped participants to perform better (i.e., selecting the correct node) without the SA in the second trial (based on their experience with the SA in the first trial), compared to the Incongruent condition. However, the ANOVA tests showed that these results were not significant.

Subjective Measures

A one-way ANOVA test revealed that mean scores for level of psychological engagement with the content explored were significantly different between the three conditions. An LSD test revealed that participants in the Congruent condition ($M = 4.26, SD = 0.37$) reported a significantly higher level of engagement than participants in the Incongruent ($M = 3.70, SD = 0.70$) and Control ($M = 3.73, SD = 0.70$) conditions (see Figure 11.9(d)); $F(2, 43) = 3.747, p = 0.032$. Additionally, significant differences were also found between the conditions on ecological validity/naturalness (i.e., congruence with real-world experience) [$F(2, 43) = 3.903, p = 0.028$], with participants in the Congruent condition ($M = 3.40, SD = 0.69$) reporting higher mean scores than participants in Control condition ($M = 2.62, SD = 0.73$), but not than the Incongruent condition ($M = 2.90, SD = 0.87$) (see Figure 11.9(d)).

No significant differences among the three conditions were found for spatial presence, negative effects, enjoyment/interest, or value/usefulness.

11.4.5 Discussion

In the Implicit Interaction Evaluation, we wanted to validate the usage of a Sentient Agent, which inferred the users' cognitive states (from physiological signals obtained through sensors worn by the users) and consequently adapted the presentation of a large dataset in order to optimize the user experience. The last experiment revealed some benefits gained from a real-time implicit interaction with a complex dataset in an immersive virtual environment. Results indicated that, when performing a task of medium complexity, participants who received SA interventions when most needed (i.e., Congruent) performed the task with a fewer number of errors and in less time than participants who received SA interventions when least needed (i.e., Incongruent), or those who did not receive any help from the SA (i.e., Control). Most importantly, 86 % of the time, the SA helped users in selecting the correct node. Further to this, when looking exclusively at the medium level of network difficulty, the SA helped users in selecting the correct node 97 % of the times of the Congruent

condition, compared to 84 % of the Incongruent and 80 % of the Control. These findings suggest that SA interventions are most effective when they support tasks of medium complexity and that the fewer or additional interventions at lower or higher levels of complexity, might impair rather than facilitate performance. Analysis of the comparison between first and second trials revealed that, regardless of their experience with the SA in the first trial, relative to participants in the Incongruent condition, Congruent SA interventions did not help participants perform better in the second trial without the SA support. Future research with more extensive measures is necessary to determine whether these SA interventions have an influence on memory retrieval. Additionally, in order to test for an interaction effect between difficulty level and condition on participants' performance, we recommend future studies use a larger sample size or adjust the complexity of the network to make the task more difficult (and participants more prone to errors) and thus increase the number of SA interventions.

Interestingly, the efficiency of the SA was also reflected in the subjective responses to the questionnaires administered after the experiment. Participants who received SA interventions when most needed reported the activity to be more engaging and perceived the content as more natural and congruent with a real-world experience (i.e., a sense of realism). Findings from this study are consistent with similar research on subliminal cueing in virtual environments previously conducted (DeVaul et al., 2003).

11.5 Conclusions

In recent years, the amount of data collected and stored has highlighted the need for effective systems that allow for their exploration in meaningful ways to generate insight and value. While computational approaches to process these large amounts of information are commonplace in many industries, human intervention is still a necessary constituent in the knowledge extraction process. To aid in this, different aspects of the human-

computer interaction can be optimized, including the characteristics of the interface itself and the input methods employed.

Intending to aid in the exploration of such large and complex datasets, we developed a system that allows for multimodal interaction in an immersive environment, exploiting the users' implicit feedback (inferring arousal and cognitive workload). Here, we presented an experimental evaluation of the system to investigate its potential benefits. To do so, we performed a series of three studies focused on the three key features that define our approach to improve data exploration: immersion, explicit interaction, and implicit interaction.

The results show that the system built here indeed enhances the interaction process in several meaningful ways. First, immersion was found to promote better recall and recognition performance and to reduce side-effects induced by exploring large data in a limited visual space. In particular, our results show that more immersive display methods, such as the use of large projection screens, are advantageous in comparison to more traditional desktop displays. Next, explicit interaction was found to be a crucial factor for stimulating a sense of presence, which led participants to become more involved in their interaction with the content explored. Specifically, allowing users to freely and actively interact with the data presented allowed them to have a more rewarding experience, in comparison to more passive exposures to the same data. Finally, the implicit interaction provided in real-time by the assistive agent provided for a more engaging and enjoyable experience, and significantly supported participants in completing tasks of medium difficulty with efficiency and efficacy. Particularly, we showed that the exploration procedure can be enhanced by dynamically providing assistance based on user states estimated from physiological signals.

Altogether, our results suggest that it is possible to substantially improve task performance through the use of automatic and adaptive interventions cued by the user's actions and mental states in an immersive environment. In order to be effective, however, an assistive agent must intervene at the right time (when most needed by users) and for a reasonable number of times to convey a seamless experience. The result is one

of a symbiotic relationship between human and computer, whereby if the user is finding the experience stressful, the agent may simplify the content to reduce cognitive workload; conversely, if the user is highly aroused, it may change the navigation speed and produce different sounds on different sides of the environment to provide greater volumes of implicit information. Synthetic environments, such as the one developed within this project, not only allow the possibility of studying automated influences in more realistic contexts which may be used by intelligent computer systems to prevent users from becoming overloaded with information but have also proven to be an excellent and effective tool in enhancing the user experience.

Acknowledgments The authors would like to thank Andrea Miotto, Ryszard Cetnarski, Alberto Betella, Rhiannon Phillips, Eva Ferrari, and Jane Lessiter for their contributions to the research reported in this chapter.

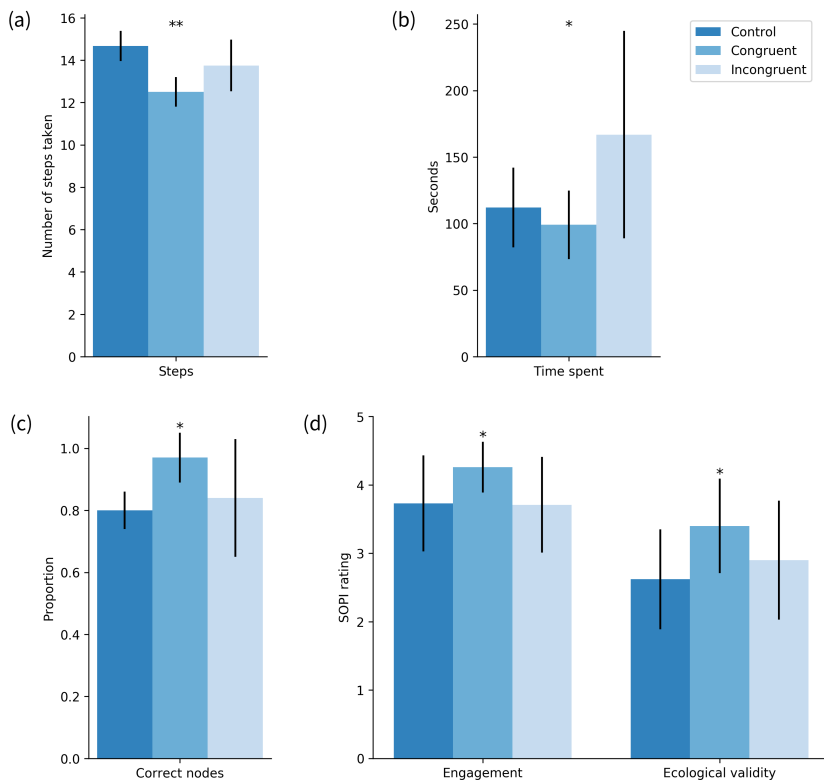


Figure 11.9: Differences between the Control, Congruent, and Incongruent conditions in the Implicit Interaction Evaluation. Objective measures correspond to the medium level of difficulty. **(a)** Participants took a significantly lower amount of steps to reach the target node in the first trial in the Medium difficulty level when the SA provided them assistance when most needed (Congruent), in comparison to when no assistance was provided (Control) or when it was provided when least needed (Incongruent). **(b)** Participants spent significantly less time to reach the target node in the first trial in the Congruent condition than in the Incongruent condition. **(c)** Participants in the Congruent condition selected a higher proportion of correct nodes than those in the Control and Incongruent conditions. **(d)** Participants reported higher levels of Engagement and Ecological validity in the Congruent condition.

Part IV

**General Discussion and
Conclusions**

Chapter 12

GENERAL DISCUSSION

Throughout this thesis, we have presented a variety of work conducted with the overall aim of advancing the development of interactive systems that are able to interact with users not only based on their conventional explicit input but also by understanding and modulating implicit cognitive and affective states: empathic systems. To achieve this, first, we asked what is needed from a technological standpoint to obtain the necessary user information. In response to this, we built a flexible sensing architecture. Next, we asked how we can infer implicit internal states. We developed this over a wide range of studies, with more specific questions for each of the methods presented. Finally, we asked how an interactive and adaptive system should be implemented and used. We answered this by working on two different systems developed to successfully assist users. In this chapter, we will discuss the insights generated through the work presented in the previous chapters. To conclude this thesis, we will provide general conclusions in the following chapter.

We started by presenting the technical architecture required to capture various user signals, process them, and react to them in Part I. Chapter 2 presented the general architecture, implemented in an immersive environment: the distributed sensing architecture (DSA). The DSA follows a layered, modular, and distributed design, allowing it to be flexibly used with a wide variety of physiological sensors, adapted to the needs of each

setting. This architecture has the capability to acquire diverse user signals and transmit them, handling synchronization to make them available for online processing to extract features from them and infer higher-level user states, as well as for the recording of the signals. Then, this information can be used by different applications to trigger events, such as adapting the interface or the content that is presented to users. Thus, the DSA serves as the basis for an empathic system, which needs to be able to collect this implicit information and react to it.

Over the next two chapters, we presented two different variations of this architecture. Chapter 3 presents a use case oriented towards the simulation of sensory overstimulation in neurodiverse individuals, including Autism Spectrum Disorder. To do this, we developed an experience based on virtual reality through a head-mounted display, coupled with a physiological sensor to acquire the electrodermal activity and heart rate (using PPG) of the users. Here, we presented the specific architecture implemented for this, as a variation of the one presented in Chapter 2. This serves as a showcase of the flexibility of the architecture, being adaptable for a virtual reality experience targeting neurodiverse phenomenology.

Next, in Chapter 4, we presented an expanded version of the architecture for usage in neurorehabilitation after stroke in serious gaming scenarios: the Socially Cooperative Cognitive Architecture for Rehabilitation (SoCCA). The SoCCA extends the DSA by integrating a more comprehensive user model: the Digital Twin. The SoCCA is designed to dynamically generate and update this Digital Twin, collecting a series of metrics, including both state and online variables, that are then analyzed to extract higher-level features, i.e., internal user states, such as stress, physical (cognitive and physical), cognitive load, and attention. Furthermore, the SoCCA integrates two new categories of devices: a functional electrical stimulation (FES) system and an exoskeleton. Both of these devices are to be modulated by the SoCCA to adapt to the patients' needs over time. Overall, this work highlights the potential of such an architecture for the implicit understanding of human cognitive and affective states, and its diversity to be adapted for a variety of use cases, including digital health, and using a variety of sensing devices.

In Part II, we introduced a series of novel methods to infer different cognitive and affective states from participants using diverse sources of information. We started this with Chapter 5, in which we presented two sequential studies aimed at inducing the unconscious processing of emotional stimuli in an immersive mixed-reality environment. Here, we used the sensing architecture previously presented in Chapter 2 (DSA) to collect a variety of user data, including electrodermal activity, pupil dilation, and EEG signals. Although the induced unconscious stimulation did not benefit the participants' performance when navigating a virtual maze, we found multiple physiological indicators of unconscious processing. This offered insights into the cognitive processes that guide the processing of subliminal emotional stimuli. Furthermore, this study served as an experimental evaluation of the DSA.

In Chapter 6, we presented a study in which we show that keystroke dynamics can be used to infer the affective state of users. In particular, we showed that two specific features of keystroke dynamics, flight times and dwell times, correlate separately with arousal and valence, respectively. This implies that affective information can be decoded from the way in which people type, without needing to analyze the content of the text. Therefore, this method can be employed without the need for specific devices, with a standard keyboard, and while preserving the privacy of the users, unlike most of the existing text-based methods.

Next, in Chapter 7 we presented a study conducted at the height of the COVID-19 pandemic in the Spring of 2020. Here, we used affective ratings to study the impact that the quarantine lockdown was having on the mental health of the general population. In particular, we found that, during this period, participants were rating neutral images more negatively than in the past, based on normative affective ratings. We interpreted this as the result of a more negative state of mind, biasing the perception of the stimuli. Furthermore, we identified different aspects of the participants' individual experiences that were associated with this negative bias. With this study, we used the novel techniques at our disposal to provide timely insights. Moreover, we also demonstrated the potential for affective ratings to be used as a valuable tool to monitor the mental health and

well-being of the general population using this implicit method, which can be delivered with more ease than traditional methods relying on questionnaires.

We continued analyzing the impact of the COVID-19 lockdown on mental health in Chapter 8, with another study that expanded on the methods of Chapter 7. Here, we targeted a larger sample size and we included additional measures to obtain a more comprehensive model of each participant. In particular, in addition to the affective ratings, we collected mouse movements, as well as text input to analyze the keystroke dynamics and perform text sentiment analysis. We also employed a questionnaire for personality and another for depression severity. Our results support our findings presented in Chapter 7 and further identify the impact of the pandemic lockdown on different demographic groups regarding age and gender. Moreover, we show the capabilities of the implicit measures that we collected in identifying variations in mental well-being, thus reaffirming their potential usage in a new generation of digital health applications.

After this, based on the insights obtained in the previous two chapters, in Chapter 9 we introduced a new method to provide affective ratings through swiping. In this method, targeted for mobile devices, stimuli are swiping either right or left, to indicate like or dislike, respectively. We showed that this method allows us to infer continuous valence information from the binary ratings provided by the participants. Furthermore, other implicit metrics (response time, swipe velocity) also are indicative of the absolute valence polarity of the stimuli. Additionally, our results also support a growing body of evidence questioning the orthogonality of arousal and valence, hinting at arousal being related to the bipolar intensity of valence. Altogether, this chapter proves the potential of this method for binary affective ratings through swiping to be used in future studies and tools. In particular, together with the results presented in the preceding chapters, we see an opportunity for this method to be used as a way to monitor fluctuations in mental well-being using a smartphone app.

Finally, in Part III, we showcased two different systems that are able to dynamically provide assistance to users when needed, as examples of the potential impact that an empathic system could have. In Chapter 10,

we presented an assistive system aimed at transferring domain knowledge. To achieve this, the system learns from successful interactions with users in a goal-oriented task, and then it provides suggestions to other users when they might benefit from them. This method, albeit simple, highlights the possibilities for providing enhanced and context-aware assistance to users, especially if augmented by additional measures for a deeper understanding of the users' needs.

Lastly, in Chapter 11 we presented the experimental evaluation of an immersive, interactive, and confluent system. We divided this study into three successive evaluations of core aspects of a data exploration experience: immersion, explicit interaction, and implicit interaction. We found that the increased immersion afforded by larger displays, and allowing users to explicitly interact with the experience were beneficial aspects. Then, we showed that interventions based on the users' unconscious reactions to the presented data, based on physiological signals, helped them in the completion of navigational tasks with improved efficiency and efficacy. This serves as the final piece of evidence within this collection of work, highlighting once again the extensive potential of empathic systems to assist users in ways not possible through conventional interaction methods.

Chapter 13

GENERAL CONCLUSIONS

Emotions are a fundamental part of human life. Historically, they have been considered by Western philosophy as separate or even opposite from rationality (Schmitter, 2016). However, thanks to profound advancements in our understanding of emotions, we now know that they play significant roles in many cognitive processes, such as perception, learning, decision-making, creativity, and memory (Tyng et al., 2017; Barrett et al., 2016). At the same time, our computing devices have been gradually becoming more ubiquitous and powerful. Nowadays, especially in high-income countries, we are increasingly surrounded by advanced devices (in the form of not only computers, but also smartphones, smartwatches, smart speakers, etc.), on which we rely for a wide variety of tasks, including working, socializing, playing, and managing some parts of our finances and fitness, for example. However, to this day, these devices lack meaningful emotional capabilities. They cannot understand how we feel or, in more practical terms, aspects such as our motivations, goals, and intentions. In this thesis, we explored the potential of bringing these advanced capabilities of understanding to interactive systems, in order to enable a new generation of empathic systems, capable of allowing for more natural interaction and an overall better user experience.

We started with the introduction of a new sensing architecture to collect multiple signals from users to enable the inference of cognitive and

affective states. This architecture is based on modern and open-source technologies, allowing for more flexibility, interoperability, and extensibility than preceding systems (Wagner et al., 2011). This enables the easy integration of new sensing devices, as well as new methods of processing the acquired signals to estimate different internal states from users. Furthermore, this architecture can be tailored to the specific needs of each application, with the addition or removal of different components, such as online processing or signal recording. We introduced different contexts for its deployment, including immersive systems using mixed reality and virtual reality, in use cases such as psychophysiological experimentation, simulation of neurodiverse phenomenology, and stroke neurorehabilitation. Overall, we believe that this sensing architecture allows for a diversity of usages not possible with previous approaches, thus having a potential impact on the many fields that could benefit from a better understanding of human implicit states.

Next, we presented a series of studies addressing specific methods of inferring user states from diverse sources, including physiological signals, keystroke dynamics, and affective ratings. With each of these studies, we advanced the state of the art in its respective field. We started by showing the usage of the previously introduced sensing architecture in two studies focused on the psychophysiological processing of emotional subliminal stimuli, proving its utility and gaining scientific insights. Then, we proposed a new method that allows us to infer affective information from text input using a standard keyboard by analyzing the way in which people type, and not the content of what they type, thus preserving privacy and obtaining valuable information without the need for special devices or actions from the users. Additionally, we conducted two timely studies at the height of the COVID-19 pandemic in the Spring of 2020, obtaining relevant indicators of the impact that the lockdown quarantine was having on the mental health and well-being of the general population, as well as showing the potential to use affective ratings to obtain such metrics. Using these insights, we then proposed a new way of providing affective ratings, advancing beyond existing tools to provide an easier and faster method through binary swiping. We believe that this method has the po-

tential to be used in clinical studies interested in long-term monitoring of mental well-being, capitalizing on fluctuating biases of emotional assessments. Furthermore, our results contribute to empirically demonstrating the relationship between arousal and valence, questioning their traditional orthogonality, as indicated by other recent studies (Kuppens et al., 2013; Haj-Ali et al., 2020). Therefore, this offers new insights into the circumplex model of affect that has been used throughout this body of work.

Finally, we presented two examples of interactive systems capable of providing assistance to their users when needed, by monitoring different metrics and reacting accordingly. The first of these, based on an online setting, learns from successful interactions with users to later provide suggestions based on different interaction factors, which could be expanded with physiological sensing in an in-person setting. The second one was an advanced navigational experience in an immersive environment, in which the system reacted to explicit and implicit interaction. These implicit factors were based on an estimation of the users' cognitive states based on physiological signals obtained through wearable devices. We showed that providing an immersive display and having both explicit and implicit interaction delivered the best results. Taken together, these two systems showcase the potential impact that can be delivered by the type of interactive and adaptive systems proposed throughout this thesis.

Overall, the work presented here offers new insights into emotion, human-computer interaction, and the intersection between both, going beyond the established field of affective computing by considering additional factors and integrating the latest knowledge in psychology and neuroscience related to the topics at hand. Moreover, this work highlights the possibilities afforded by enabling synthetic systems to understand and modulate human cognitive and affective states. These include enhancing the user experience to be more natural, with more intelligent systems that are able to proactively react to how the users are feeling to provide adaptations and assistance, thus resulting in improved efficacy and efficiency.

We believe that our results will have a general impact on fields of research such as affective science and human-computer interaction, and, particularly, on digital health. We envision a new generation of digital

health systems that are able to capture a wider range of relevant information from users while simultaneously requiring less effort, by capitalizing on implicit understanding. An example of this would be analyzing the way that patients at risk of depression type or rate affective stimuli, as shown in our studies, in order to automatically monitor the fluctuations in their mental well-being. Then, the system could intervene by offering some mitigation actions or providing professional attention by therapists. This type of novel systems for digital health could facilitate mental healthcare for both clinicians and patients alike, with a potential impact in preventing the development of negative conditions, diagnosing individuals who might be at risk of developing or worsening disorders, monitoring at-risk individuals, and intervening when support is needed.

We believe future work should be focused on two main areas. The first of these is to continue the development of closed-loop interaction systems, grounded on the latest psychological and neuroscientific theories, to further develop the usage of the inferred internal states to have an impact throughout the interactive process. This would also serve as an additional validation of some of the methods presented here. The other area of future work would include the conduction of longitudinal studies, tracking the implicit metrics presented here over time, together with additional assessments such as clinically-validated questionnaires. This would reinforce the potential utility of these features to monitor the mental state of individuals and have a positive impact on their lives, improving their health and well-being.

BIBLIOGRAPHY

- Acharya, U. R., Joseph, K. P., Kannathal, N., Lim, C. M., and Suri, J. S. (2006). Heart rate variability: A review.
- Adomavicius, G., Mobasher, B., Ricci, F., and Tuzhilin, A. (2011). What a Recommender System Knows About Contextual Factors. *Springer*, pages 67–80.
- Adomavicius, G. and Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734–749.
- Agrafioti, F., Hatzinakos, D., and Anderson, A. K. (2012). ECG pattern analysis for emotion detection. *IEEE Transactions on Affective Computing*, 3(1):102–115.
- Ahmed, S. (2017). Theory of sustained optimal challenge in teaching and learning. *Proceedings of the Human Factors and Ergonomics Society*, 2017-October:407–411.
- Akçay, M. B. and Oğuz, K. (2020). Speech emotion recognition: Emotional models, databases, features, preprocessing methods, supporting modalities, and classifiers. *Speech Communication*, 116(October 2019):56–76.
- Alfano, P. L. and Michel, G. F. (1990). Restricting the field of view:

- Perceptual and performance effects. *Perceptual and Motor Skills*, 70(1):35–45.
- Ali, H. (2020). Speech Emotion Recognition Survey. *Journal of Mechanics of Continua and Mathematical Sciences*, 15(9).
- Allen, J. (2007). Photoplethysmography and its application in clinical physiological measurement.
- Andre, E. (2013). Exploiting Unconscious User Signals in Multimodal Human-Computer Interaction. *ACM Transactions on Multimedia Computing*, 9:1–5.
- Aranha, R. V., Correa, C. G., and Nunes, F. L. S. (2019). Adapting software with Affective Computing: a systematic review. *IEEE Transactions on Affective Computing*, 3045(c):1–1.
- Araya León, M. J., López-Carral, H., Mura, A., and Verschure, P. F. (2023). Environmental and emotional well-being in the spatial experience in the restaurant context, from a work perspective. *Manuscript in preparation*.
- Argasiński, J. K. and Węgrzyn, P. (2019). Affective patterns in serious games. *Future Generation Computer Systems*, 92:526–538.
- Arthur, K. (1996). Effects of field of view on task performance with head-mounted displays. In *Conference on Human Factors in Computing Systems - Proceedings*, pages 29–30.
- Arya, R., Singh, J., and Kumar, A. (2021). A survey of multidisciplinary domains contributing to affective computing. *Computer Science Review*, 40:100399.
- Asabere, N. Y. (2012). A survey of personalized television and video recommender systems and techniques. *Information and Communication Technology Research*, 2(7):602–608.

- Association, A. P. et al. (2013). *Diagnostic and statistical manual of mental disorders (DSM-5®)*. American Psychiatric Pub.
- Bagby, R. M., Ryder, A. G., Schuller, D. R., and Marshall, M. B. (2004). The hamilton depression rating scale: has the gold standard become a lead weight? *American Journal of Psychiatry*, 161(12):2163–2177.
- Bai, Y., Lin, C.-C., Lin, C.-Y., Chen, J.-Y., Chue, C.-M., and Chou, P. (2004). Survey of stress reactions among health care workers involved with the sars outbreak. *Psychiatric Services*, 55(9):1055–1057.
- Ball, R. and North, C. (2005). Effects of tiled high-resolution display on basic visualization and navigation tasks. In *Conference on Human Factors in Computing Systems - Proceedings*, pages 1196–1199.
- Ballester, B. R., Lathe, A., Duarte, E., Duff, A., and Verschure, P. F. (2015). A wearable bracelet device for promoting arm use in stroke patients. In *NEUROTECHNIX*, pages 24–31.
- Ballester, B. R., Maier, M., Duff, A., Cameirão, M., Bermúdez, S., Duarte, E., Cuxart, A., Rodríguez, S., Mozo, R. M. S. S., and Verschure, P. F. (2019). A critical time window for recovery extends beyond one-year post-stroke. *Journal of Neurophysiology*, 122(1):350–357.
- Ballester, B. R., Maier, M., San Segundo Mozo, R. M., Castañeda, V., Duff, A., and M J Verschure, P. F. (2016). Counteracting learned non-use in chronic stroke patients with reinforcement-induced movement therapy. *Journal of neuroengineering and rehabilitation*, 13(1):74.
- Banerjee, D. (2020). The impact of covid-19 pandemic on elderly mental health. *International journal of geriatric psychiatry*, 35(12):1466.
- Baños, R. M., Botella, C., Alcañiz, M., Liaño, V., Guerrero, B., and Rey, B. (2004). Immersion and emotion: Their impact on the sense of presence. *Cyberpsychology and Behavior*, 7(6):734–741.

- Barbisch, D., Koenig, K. L., and Shih, F.-Y. (2015). Is there a case for quarantine? Perspectives from SARS to Ebola. *Disaster medicine and public health preparedness*, 9(5):547–553.
- Barrett, L. F. (2017a). *How emotions are made: The secret life of the brain*. Pan Macmillan.
- Barrett, L. F. (2017b). The theory of constructed emotion: an active inference account of interoception and categorization. *Social cognitive and affective neuroscience*, 12(1):1–23.
- Barrett, L. F. (2018). *How Emotions Are Made*. Pan Macmillan.
- Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., and Pollak, S. D. (2019). Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements. *Psychological Science in the Public Interest*, 20(1):1–68.
- Barrett, L. F., Lewis, M., and Haviland-Jones, J. M., editors (2016). *Handbook of Emotions*. Guilford Press, fourth edition.
- Barsy, B., Kocsis, K., Magyar, A., Babiczky, Á., Szabó, M., Veres, J. M., Hillier, D., Ulbert, I., Yizhar, O., and Mátyás, F. (2020). Associative and plastic thalamic signaling to the lateral amygdala controls fear behavior. *Nature Neuroscience*, 23(5):625–637.
- Bayliss, P. (2007). Beings in the game-world: Characters, avatars, and players. In *Proceedings - 4th Australasian Conference on Interactive Entertainment, IE 2007*.
- Beatty, J. and Lucero-Wagoner, B. (2000). The pupillary system. In *Handbook of psychophysiology (2nd ed.)*, pages 142–162. Cambridge University Press.
- Bech, P., Allerup, P., Reisby, N., and Gram, L. (1984). Assessment of symptom change from improvement curves on the hamilton depression scale in trials with antidepressants. *Psychopharmacology*, 84(2):276–281.

- Beck, A., Steer, R., and Brown, G. (1996). Manual for the Beck depression inventory-II. *San Antonio, TX: Psychological Corporation.*
- Benedek, M. and Kaernbach, C. (2010). A continuous measure of phasic electrodermal activity. *Journal of Neuroscience Methods*, 190(1):80–91.
- Berenbaum, H. and Oltmanns, T. F. (1992). Emotional experience and expression in schizophrenia and depression. *Journal of abnormal psychology*, 101(1):37.
- Bergadano, F., Gunetti, D., and Picardi, C. (2003). User authentication through keystroke dynamics. *ACM Transactions on Information and System Security*, 5(4):367–397.
- Bernardet, U., Badia, S. B. i., Duff, A., Inderbitzin, M., Groux, S. L., Manzolli, J., Mathews, Z., Mura, A., Våljamäe, A., and Verschure, P. F. M. J. (2010). The eXperience Induction Machine: A New Paradigm for Mixed-Reality Interaction Design and Psychological Experimentation. In Dubois, E., Gray, P., and Nigay, L., editors, *The Engineering of Mixed Reality Systems*, Human-Computer Interaction Series, pages 357–379. Springer London, London.
- Bernardet, U., Våljamäe, A., Inderbitzin, M., Wierenga, S., Mura, A., and Verschure, P. F. M. J. (2011). Quantifying human subjective experience and social interaction using the eXperience Induction Machine. *Brain Research Bulletin*, 85(5):305–312.
- Bernat, E., Bunce, S., and Shevrin, H. (2001). Event-related brain potentials differentiate positive and negative mood adjectives during both supraliminal and subliminal visual processing. *International Journal of Psychophysiology*, 42(1):11–34.
- Berridge, K. C. (2019). Affective valence in the brain: modules or modes? *Nature Reviews Neuroscience*.

- Betella, A., Carvalho, R., Sanchez-Palencia, J., Bernardet, U., and Verschure, P. F. M. J. (2012). Embodied interaction with complex neuronal data in mixed-reality. *Virtual Reality International Conference (VRIC 2012)*, page 3.
- Betella, A., Cetnarski, R., Zucca, R., Arsiwalla, X. D., Martinez, E., Omedas, P., Mura, A., and Verschure, P. F. M. J. (2014a). BrainX 3: embodied exploration of neural data. In *Proceedings of the 2014 Virtual Reality International Conference*, page 37. ACM.
- Betella, A., Pacheco, D., Zucca, R., Arsiwalla, X. D., Omedas, P., Lanatà, A., Mazzei, D., Tognetti, A., Greco, A., Carbonaro, N., Wagner, J., Lingenfelser, F., André, E., Rossi, D. D., and Verschure, P. F. (2014b). Interpreting Psychophysiological States Using Unobtrusive Wearable Sensors in Virtual Reality. *ACHI 2014, The Seventh International Conference on Advances in Computer-Human Interactions*, pages 331–336.
- Betella, A. and Verschure, P. F. M. J. (2016). The affective slider: A digital self-assessment scale for the measurement of human emotions. *PLoS ONE*, 11(2):1–11.
- Betella, A., Zucca, R., Cetnarski, R., Greco, A., Lanatà, A., Mazzei, D., Tognetti, A., Arsiwalla, X. D., Omedas, P., De Rossi, D., and Verschure, P. F. (2014c). Inference of human affective states from psychophysiological measurements extracted under ecologically valid conditions. *Frontiers in Neuroscience*, 8(SEP):66.
- Beyeler, A., Namburi, P., Glober, G. F., Simonnet, C., Calhoun, G. G., Conyers, G. F., Luck, R., Wildes, C. P., and Tye, K. M. (2016). Divergent routing of positive and negative information from the amygdala during memory retrieval. *Neuron*, 90(2):348–361.
- Bo, H.-X., Li, W., Yang, Y., Wang, Y., Zhang, Q., Cheung, T., Wu, X., and Xiang, Y.-T. (2020). Posttraumatic stress symptoms and attitude toward crisis mental health services among clinically stable patients with covid-19 in china. *Psychological medicine*, pages 1–2.

- Boag, S. (2009). Making Sense of Subliminal Perception. In Columbus, A. M., editor, *Advances in Psychology Research*, volume 54, pages 117–139. New York: Nova Science Publishers.
- Bontchev, B. (2016). Adaptation in affective video games: A literature review. *Cybernetics and Information Technologies*, 16(3):3–34.
- Bornstein, R. F. and Pittman, T. S. (1992). *Perception Without Awareness: Cognitive, Clinical, and Social Perspectives*. Guilford Press, New York, NY, US.
- Bortolla, R., Cavicchioli, M., Galli, M., Verschure, P. F. M. J., and Maffei, C. (2019). A comprehensive evaluation of emotional responsiveness in borderline personality disorder: a support for hypersensitivity hypothesis. *Borderline Personality Disorder and Emotion Dysregulation*, 6(1).
- Bos, D. O. (2006). EEG-based Emotion Recognition. *The influence of Visual and Auditory Stimuli*, 56(3):1–17.
- Bowman, D. A. and McMahan, R. P. (2007). Virtual reality: How much immersion is enough? *Computer*, 40(7):36–43.
- Bradley, M. M. and Lang, P. J. (1994). Measuring emotion: The self-assessment manikin and the semantic differential. *Journal of Behavior Therapy and Experimental Psychiatry*, 25(1):49–59.
- Bradley, M. M., Miccoli, L., Escrig, M. A., and Lang, P. J. (2008). The pupil as a measure of emotional arousal and autonomic activation. *Psychophysiology*, 45(4):602–607.
- Braun, H. I., Jackson, D. N., and Wiley, D. E. (2001). *The role of constructs in psychological and educational measurement*. Routledge.
- Brooke, A. H. and Harrison, N. A. (2016). Neuroimaging and Emotion. In *Stress: Concepts, Cognition, Emotion, and Behavior: Handbook of Stress*, pages 251–259. Elsevier Inc.

- Brooks, S. K., Webster, R. K., Smith, L. E., Woodland, L., Wessely, S., Greenberg, N., and Rubin, G. J. (2020). The psychological impact of quarantine and how to reduce it: rapid review of the evidence. *The Lancet*, 395(10227):912–920.
- Brusilovsky, P. (2007). Adaptive navigation support. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 4321 LNCS:263–290.
- Bruynseels, K., Santoni de Sio, F., and van den Hoven, J. (2018). Digital twins in health care: ethical implications of an emerging engineering paradigm. *Frontiers in genetics*, 9:31.
- Buder, J. and Schwind, C. (2012). Learning with personalized recommender systems: A psychological view. *Computers in Human Behavior*, 28(1):207–216.
- Burke, R. (2002). Hybrid recommender systems: Survey and experiments. *User modeling and user-adapted interaction*, 12(4):331–370.
- Burkhardt, F. and Sendlmeier, W. F. (2000). Verification of acoustical correlates of emotional speech using formant-synthesis. In *Proc. ISCA Workshop on Speech and Emotion*, pages 151–156.
- Cacioppo, J. T., Hawkley, L. C., Crawford, L. E., Ernst, J. M., Burleson, M. H., Kowalewski, R. B., Malarkey, W. B., Van Cauter, E., and Berntson, G. G. (2002). Loneliness and health: Potential mechanisms. *Psychosomatic Medicine*, 64(3):407–417.
- Cacioppo, J. T., Hawkley, L. C., and Thisted, R. A. (2010). Perceived social isolation makes me sad: 5-year cross-lagged analyses of loneliness and depressive symptomatology in the chicago health, aging, and social relations study. *Psychology and Aging*, 25(2):453–463.
- Cacioppo, J. T., Hughes, M. E., Waite, L. J., Hawkley, L. C., and Thisted, R. A. (2006). Loneliness as a specific risk factor for depressive symptoms: Cross-sectional and longitudinal analyses. *Psychology and Aging*, 21(1):140–151.

- Calvo, R. A., D’Mello, S., Gratch, J. M., and Kappas, A., editors (2015). *The Oxford Handbook of Affective Computing*. Oxford Library of Psychology.
- Cannon, W. B. (1927). The james-lange theory of emotions: A critical examination and an alternative theory. *The American Journal of Psychology*, 39(1/4):106–124.
- Carbonaro, N., Greco, A., Anania, G., Dalle Mura, G., Tognetti, A., Scilingo, E. P., De Rossi, D., and Lanatà, A. (2012). Unobtrusive Physiological and Gesture Wearable Acquisition System: A Preliminary Study on Behavioral and Emotional Correlations. In *GLOBAL HEALTH 2012, The First International Conference on Global Health Challenges*, pages 88–92.
- Caruelle, D., Shams, P., Gustafsson, A., and Lervik-Olsen, L. (2022). Affective computing in marketing: practical implications and research opportunities afforded by emotionally intelligent machines. *Marketing Letters*, 33(1):163–169.
- Casasanto, D. (2009). Embodiment of Abstract Concepts: Good and Bad in Right- and Left-Handers. *Journal of Experimental Psychology: General*, 138(3):351–367.
- Caspi, A., Harrington, H. L., Moffitt, T. E., Milne, B. J., and Poulton, R. (2006). Socially isolated children 20 years later: Risk of cardiovascular disease. *Archives of Pediatrics and Adolescent Medicine*, 160(8):805–811.
- Centers for Disease Control and Prevention (2017). Quarantine and isolation. <https://www.cdc.gov/quarantine/index.html>.
- Cervera-Torres, S., Ruiz Fernández, S., Lachmair, M., Riekert, M., and Gerjets, P. (2021). Altering emotions near the hand: Approach-avoidance swipe interactions modulate the perceived valence of emotional pictures. *Emotion*, 21(1):220–225.

- Cetnarski, R., Betella, A., Miotto, A., Zucca, R., Arsiwalla, X. D., Omedas, P., Freeman, J., and Verschure, P. F. M. J. (2015). Symbiotic Adaptive Interfaces: A Case Study Using BrainX3. In *Symbiotic Interaction*, pages 33–44. Springer.
- Cetnarski, R., Betella, A., Prins, H., Kouider, S., and Verschure, P. F. M. J. (2014). Subliminal Response Priming in Mixed Reality: The Ecological Validity of a Classic Paradigm of Perception. *Presence: Teleoperators and Virtual Environments*.
- Chambon, V., Farrer, C., Pacherie, E., Jacquet, P. O., Leboyer, M., and Zalla, T. (2017). Reduced sensitivity to social priors during action prediction in adults with autism spectrum disorders. *Cognition*, 160:17–26.
- Charles, D., Holmes, D., Charles, T., and McDonough, S. (2020). Virtual Reality Design for Stroke Rehabilitation. *Advances in Experimental Medicine and Biology*, 1235(11):53–87.
- Cheesman, J. and Merikle, P. M. (1984). Priming with and without awareness. *Perception & Psychophysics*, 36(4):387–395.
- Cifrek, M., Medved, V., Tonković, S., and Ostojić, S. (2009). Surface emg based muscle fatigue evaluation in biomechanics. *Clinical biomechanics*, 24(4):327–340.
- Ciuk, D., Troy, A., and Jones, M. (2015). Measuring emotion: Self-reports vs. physiological indicators. *Physiological Indicators (April 16, 2015)*.
- Clark, L. A. and Watson, D. (1991). Tripartite model of anxiety and depression: psychometric evidence and taxonomic implications. *Journal of abnormal psychology*, 100(3):316.
- Coan, J. A. and Allen, J. J. (2004). Frontal EEG asymmetry as a moderator and mediator of emotion. *Biological Psychology*, 67(1-2):7–50.

- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Conti, A. A. (2020). Historical and methodological highlights of quarantine measures: From ancient plague epidemics to current coronavirus disease (COVID-19) pandemic. *Acta Biomedica*, 91(2):226–229.
- Cornwell, E. Y. and Waite, L. J. (2009). Social disconnectedness, perceived isolation, and health among older adults. *Journal of Health and Social Behavior*, 50(1):31–48.
- Cosoli, G., Poli, A., Scalise, L., and Spinsante, S. (2021). Measurement of multimodal physiological signals for stimulation detection by wearable devices. *Measurement: Journal of the International Measurement Confederation*, 184(April):109966.
- Courtney, C. G., Dawson, M. E., Schell, A. M., Iyer, A., and Parsons, T. D. (2010). Better than the real thing: Eliciting fear with moving and static computer-generated stimuli. *International Journal of Psychophysiology*, 78(2):107–114.
- Cowley, B., Filetti, M., Lukander, K., Torniainen, J., Henelius, A., Ahonen, L., Barral, O., Kosunen, I., Valtonen, T., Huotilainen, M., Ravaja, N., and Jacucci, G. (2016). The Psychophysiology Primer: A Guide to Methods and a Broad Review with a Focus on Human-Computer Interaction. *Foundations and Trends in Human-Computer Interaction*, 9(3-4):150–307.
- Cox, C. L., Uddin, L. Q., Di martino, A., Castellanos, F. X., Milham, M. P., and Kelly, C. (2012). The balance between feeling and knowing: Affective and cognitive empathy are reflected in the brain’s intrinsic functional dynamics. *Social Cognitive and Affective Neuroscience*, 7(6):727–737.
- Crick, T. (2011). The game body: Toward a phenomenology of contemporary video gaming. *Games and Culture*, 6(3):259–269.

- Critchley, H. D. (2002). Review: Electrodermal Responses: What Happens in the Brain. *The Neuroscientist*, 8(2):132–142.
- Croatti, A., Gabellini, M., Montagna, S., and Ricci, A. (2020). On the integration of agents and digital twins in healthcare. *Journal of Medical Systems*, 44(9):1–8.
- da Silva Cameirão, M., Bermúdez I Badía, S., Duarte, E., and Verschure, P. F. M. J. (2011). Virtual reality based rehabilitation speeds up functional recovery of the upper extremities after stroke: a randomized controlled pilot study in the acute phase of stroke using the rehabilitation gaming system. *Restorative neurology and neuroscience*, 29(5):287–298.
- Dal Santo, T., Sun, Y., Wu, Y., He, C., Wang, Y., Jiang, X., Li, K., Bonardi, O., Krishnan, A., Boruff, J. T., et al. (2022). Systematic review of mental health symptom changes by sex or gender in early-covid-19 compared to pre-pandemic. *Scientific Reports*, 12(1):1–14.
- Damasio, A. and Carvalho, G. B. (2013). The nature of feelings: Evolutionary and neurobiological origins. *Nature Reviews Neuroscience*, 14(2):143–152.
- Damasio, A. R. (1994). *Descartes' Error: Emotion, Reason and the Human Brain*. Putnam.
- Damasio, A. R. (1996). The somatic marker hypothesis and the possible functions of the prefrontal cortex. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 351(1346):1413–1420.
- Damian, M. F. (2001). Congruity effects evoked by subliminally presented primes: Automaticity rather than semantic processing. *Journal of Experimental Psychology: Human Perception and Performance*, 27(1):154–165.

- Dan-Glauser, E. S. and Scherer, K. R. (2011). The Geneva affective picture database (GAPED): a new 730-picture database focusing on valence and normative significance. *Behavior Research Methods*.
- Darwin, C. (1872). *The Expression of the Emotions in Man and Animals*. John Murray.
- David, G. and Cambre, C. (2016). Screened Intimacies: Tinder and the Swipe Logic. *Social Media and Society*, 2(2).
- Davis, M. H. (2022). Empathy. In *Reference Module in Neuroscience and Biobehavioral Psychology*. Elsevier.
- Delorme, A. and Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, 134(1):9–21.
- Demain, S., Wiles, R., Roberts, L., and McPherson, K. (2006). Recovery plateau following stroke: Fact or fiction? *Disability and Rehabilitation*, 28(13-14):815–821.
- Demyttenaere, K., Kiekens, G., Bruffaerts, R., Mortier, P., Gorwood, P., Martin, L., and Di Giannantonio, M. (2020). Outcome in depression (ii): beyond the hamilton depression rating scale. *CNS spectrums*, pages 1–22.
- Detenber, B. H., Simons, R. F., and Bennett, G. G. (1998). Roll 'em!: The effects of picture motion on emotional responses. *Journal of Broadcasting and Electronic Media*, 42(1):113–127.
- DeVaul, R. W., Pentland, A. S., and Corey, V. R. (2003). The memory glasses: Subliminal vs.overt memory support with imperfect information. In *Proceedings - International Symposium on Wearable Computers, ISWC*, pages 146–153.

- Diemer, J., Alpers, G. W., Peperkorn, H. M., Shiban, Y., and Mühlberger, A. (2015). The impact of perception and presence on emotional reactions: A review of research in virtual reality. *Frontiers in Psychology*, 6(JAN):1–9.
- D’Mello, S. and Calvo, R. A. (2013). Beyond the Basic Emotions: What Should Affective Computing Compute? *Conference on Human Factors in Computing Systems - Proceedings*, 2013-April:2287–2294.
- Dodson, J. D. (1915). The relation of strength of stimulus to rapidity of habit-formation in the kitten. *Journal of Animal Behavior*, 5(4):330–336.
- Dolan, R. J. (2002). Emotion, Cognition, and Behavior. *Science*, 298(November):1191–1195.
- Douglass, A. M., Kucukdereli, H., Ponserre, M., Markovic, M., Gründemann, J., Strobel, C., Alcalá Morales, P. L., Conzelmann, K.-K., Lüthi, A., and Klein, R. (2017). Central amygdala circuits modulate food consumption through a positive-valence mechanism. *Nature neuroscience*, 20(10):1384–1394.
- Dowland, P. S. and Furnell, S. M. (2005). A Long-Term Trial of Keystroke Profiling Using Digraph, Trigraph and Keyword Latencies. In *Security and Protection in Information Processing Systems*, pages 275–289.
- Dunn, B. D., Dalgleish, T., Lawrence, A. D., Cusack, R., and Ogilvie, A. D. (2004). Categorical and dimensional reports of experienced affect to emotion-inducing pictures in depression. *Journal of abnormal psychology*, 113(4):654.
- Dzedzickis, A., Kaklauskas, A., and Bucinskas, V. (2020). Human emotion recognition: Review of sensors and methods. *Sensors*, 20(3).
- Ehrsson, H. H. (2012). The Concept of Body Ownership and Its Relation to Multisensory Integration. In *The New Handbook of Multisensory Process*, pages 775–792. MIT press Cambridge, MA.

- Eimer, M. and Schlaghecken, F. (2003). Response facilitation and inhibition in subliminal priming. *Biological Psychology*, 64(1-2):7–26.
- Ekkekakis, P. (2013). *The measurement of affect, mood, and emotion: A guide for health-behavioral research*. Cambridge University Press.
- Ekman, P. (1992). An argument for basic emotions.
- Ekman, P. and Cordaro, D. (2011). What is meant by calling emotions basic. *Emotion Review*, 3(4):364–370.
- Ekman, P. and Friesen, W. V. (1976). *Pictures of Facial Affect*. Consulting Psychologists Press, Palo Alto, CA.
- Ekman, P. and Oster, H. (1979). Facial Expressions of Emotion. *Annual Review of Psychology*, 30(1):527–554.
- Elliott, R., Bohart, A. C., Watson, J. C., and Greenberg, L. S. (2011). Empathy. *Psychotherapy*, 48(1):43–49.
- Epp, C., Lippold, M., and Mandryk, R. L. (2011). Identifying emotional states using keystroke dynamics. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*, page 715.
- Fagherazzi, G., Goetzinger, C., Rashid, M. A., Aguayo, G. A., and Huiart, L. (2020). Digital health strategies to fight covid-19 worldwide: challenges, recommendations, and a call for papers. *Journal of Medical Internet Research*, 22(6):e19284.
- Fanselow, M. S. and Wassum, K. M. (2016). The origins and organization of vertebrate pavlovian conditioning. *Cold Spring Harbor perspectives in biology*, 8(1):a021717.
- Fasel, B. and Luetttin, J. (2003). Automatic facial expression analysis: A survey. *Pattern Recognition*, 36(1):259–275.

- Fernández-Aranda, F., Casas, M., Claes, L., Bryan, D. C., Favaro, A., Granero, R., Gudiol, C., Jiménez-Murcia, S., Karwautz, A., Le Grange, D., et al. (2020). Covid-19 and implications for eating disorders. *European Eating Disorders Review*, 28(3):239.
- Ferscha, A., editor (2013). *Human Computer Confluence: The Next Generation Humans and Computers Research Agenda*. Institute for Pervasive Computing, Johannes Kepler University Linz.
- Flavell, S. W., Gogolla, N., Lovett-barron, M., and Zelikowsky, M. (2022). The emergence and influence of internal states. *Neuron*, pages 1–26.
- Flowers, J. H. (2005). Thirteen Years of Reflection on Auditory Graphing: Promises , Pitfalls , and Potential New Directions. In *International Conference on Auditory Display*, pages 1–5.
- Folkman, S. and Lazarus, R. S. (1988). The relationship between coping and emotion: Implications for theory and research. *Social Science and Medicine*, 26(3):309–317.
- Fried, E. I. and Nesse, R. M. (2015). Depression sum-scores don't add up: why analyzing specific depression symptoms is essential. *BMC medicine*, 13(1):1–11.
- Frijda, N. H. et al. (1986). *The emotions*. Cambridge University Press.
- Fuzaro, A. C., Guerreiro, C. T., Galetti, F. C., Jucá, R. B. V. M., and de Araujo, J. E. (2012). Modified constraint-induced movement therapy and modified forced-use therapy for stroke patients are both effective to promote balance and gait improvements. *Revista brasileira de fisioterapia (Sao Carlos (Sao Paulo, Brazil))*, 16(2):157–165.
- Gaggioli, A., Ferscha, A., Riva, G., Dunne, S., and Viaud-Delmon, I., editors (2016). *Human Computer Confluence: Transforming Human Experience through Symbiotic Technologies*. De Gruyter.

- Ge, M. and Persia, F. (2017). A survey of multimedia recommender systems: Challenges and opportunities. *International Journal of Semantic Computing*, 11(03):411–428.
- Gervasi, R., Barravecchia, F., Mastrogiacomo, L., and Franceschini, F. (2022). Applications of affective computing in human-robot interaction: State-of-art and challenges for manufacturing. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, page 09544054221121888.
- Gibbons, R. D., Clark, D. C., and Kupfer, D. J. (1993). Exactly what does the hamilton depression rating scale measure? *Journal of psychiatric research*, 27(3):259–273.
- Girardi, D., Lanubile, F., and Novielli, N. (2017). Emotion detection using noninvasive low cost sensors. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 125–130.
- Girouard, A., Solovey, E. T., and Jacob, R. J. (2013). Designing a passive brain computer interface using real time classification of functional near-infrared spectroscopy. *International Journal of Autonomous and Adaptive Communications Systems*, 6(1):26–44.
- Goldberg, D., Nichols, D., Oki, B. M., and Terry, D. (1992). Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, 35(12):61–71.
- Goldman, S., Wang, C., Salgado, M. W., Greene, P. E., Kim, M., and Rapin, I. (2009). Motor stereotypies in children with autism and other developmental disorders. *Developmental Medicine & Child Neurology*, 51(1):30–38.
- Gomes, E., Pedroso, F. S., and Wagner, M. B. (2008). Hipersensibilidade auditiva no transtorno do espectro autístico. *Pró-Fono Revista de Atualização Científica*, 20(4):279–284.

- Grabe, M. E., Lombard, M., Reich, R. D., Bracken, C. C., and Ditton, T. B. (1999). The role of screen size in viewer experiences of media content. *Visual Communication Quarterly*, 6(2):4–9.
- Grabenhorst, F., Rolls, E. T., and Parris, B. A. (2008). From affective value to decision-making in the prefrontal cortex. *European Journal of Neuroscience*, 28(9):1930–1939.
- Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., Goj, R., Jas, M., Brooks, T., Parkkonen, L., and Hämäläinen, M. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, 7.
- Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., Parkkonen, L., and Hämäläinen, M. S. (2014). MNE software for processing MEG and EEG data. *NeuroImage*, 86:446–460.
- Grechuta, K., Ballester, B. R., Munné, R. E., Bernal, T. U., Hervás, B. M., Mohr, B., Pulvermüller, F., San Segundo, R., and Verschure, P. F. (2020). Multisensory cueing facilitates naming in aphasia. *Journal of NeuroEngineering and Rehabilitation*, 17(1):1–11.
- Greene, S., Thapliyal, H., and Caban-Holt, A. (2016a). A Survey of Affective Computing for Stress Detection: Evaluating technologies in stress detection for better health. *IEEE Consumer Electronics Magazine*, 5(4):44–56.
- Greene, S., Thapliyal, H., and Caban-Holt, A. (2016b). A survey of affective computing for stress detection: Evaluating technologies in stress detection for better health. *IEEE Consumer Electronics Magazine*, 5(4):44–56.
- Grivich, M. (2013). LSL Validation.
- Gross, J. J. and Barrett, L. F. (2011). Emotion Generation and Emotion Regulation: One or Two Depends on Your Point of View. *Emot Rev.*, 3(1):8–16.

- Grudin, J. (2017). From tool to partner: The evolution of human-computer interaction. *Synthesis Lectures on Human-Centered Interaction*, 10(1):i–183.
- Gunawardhane, S. D., De Silva, P. M., Kulathunga, D. S., and Arunatileka, S. M. (2013). Non invasive human stress detection using key stroke dynamics and pattern variations. In *International Conference on Advances in ICT for Emerging Regions, ICTer 2013 - Conference Proceedings*, pages 240–247. IEEE.
- Gunes, H., Schuller, B., Pantic, M., and Cowie, R. (2011). Emotion representation, analysis and synthesis in continuous space: A survey. *2011 IEEE International Conference on Automatic Face and Gesture Recognition and Workshops, FG 2011*, pages 827–834.
- Guo, H. W., Huang, Y. S., Lin, C. H., Chien, J. C., Haraikawa, K., and Shieh, J. S. (2016). Heart Rate Variability Signal Features for Emotion Recognition by Using Principal Component Analysis and Support Vectors Machine. *Proceedings - 2016 IEEE 16th International Conference on Bioinformatics and Bioengineering, BIBE 2016*, pages 274–277.
- Guo, P. J., Kim, J., and Rubin, R. (2014). How video production affects student engagement: An empirical study of MOOC videos. In *L@S 2014 - Proceedings of the 1st ACM Conference on Learning at Scale*, pages 41–50.
- Gupta, V., Hanges, P. J., and Dorfman, P. (2002). Cultural clusters: Methodology and findings. *Journal of World Business*, 37(1):11–15.
- Hagen, M. A., Jones, R. K., and Reed, E. S. (1978). On a neglected variable in theories of pictorial perception: Truncation of the visual field. *Perception & Psychophysics*, 23(4):326–330.
- Haj-Ali, H., Anderson, A. K., and Kron, A. (2020). Comparing three models of arousal in the human brain. *Social Cognitive and Affective Neuroscience*, 15(1):1–11.

- Hamilton, M. (1960). A rating scale for depression. *Journal of neurology, neurosurgery, and psychiatry*, 23(1):56.
- Haney, C. (2003). Mental health issues in long-term solitary and “super-max” confinement. *Crime & Delinquency*, 49(1):124–156.
- Harmon-Jones, E. (2003). Clarifying the emotive functions of asymmetrical frontal cortical activity. *Psychophysiology*, 40(6):838–848.
- Hasan, M. S. and Yu, H. (2017). Innovative developments in hci and future trends. *International Journal of Automation and Computing*, 14(1):10–20.
- Hawkey, L. C. and Capitanio, J. P. (2015). Perceived social isolation, evolutionary fitness and health outcomes: A lifespan approach. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1669).
- Haywood, N. and Cairns, P. (2006). Engagement with an interactive museum exhibit. In *People and Computers XIX - The Bigger Picture, Proceedings of HCI 2005*, pages 113–129.
- Hendrix, C. and Barfield, W. (1996). Presence within virtual environments as a function of visual display parameters. *Presence: Teleoperators and Virtual Environments*, 5(3):274–289.
- Henelius, A. and Torniaainen, J. (2018). MIDAS: Open-source framework for distributed online analysis of data streams. *SoftwareX*, 7:156–161.
- Hidaka, Y., Han, C. E., Wolf, S. L., Winstein, C. J., and Schweighofer, N. (2012). Use It and Improve It or Lose It : Interactions between Arm Function and Use in Humans Post-stroke. *PLoS Computational Biology*, 8(2).
- Holmes, E. A., O’Connor, R. C., Perry, V. H., Tracey, I., Wessely, S., Arseneault, L., Ballard, C., Christensen, H., Cohen Silver, R., Overall,

- I., Ford, T., John, A., Kabir, T., King, K., Madan, I., Michie, S., Przybylski, A. K., Shafran, R., Sweeney, A., Worthman, C. M., Yardley, L., Cowan, K., Cope, C., Hotopf, M., and Bullmore, E. (2020). Multi-disciplinary research priorities for the COVID-19 pandemic: a call for action for mental health science. *The Lancet Psychiatry*, 7(6):547–560.
- Holt-Lunstad, J. (2017). The Potential Public Health Relevance of Social Isolation and Loneliness: Prevalence, Epidemiology, and Risk Factors. *Public Policy & Aging Report*, 27(4):127–130.
- Hommel, B. (2000). The prepared reflex: automaticity and control in stimulus-response translation. *Control of cognitive processes: attention and performance XVIII*, pages 247–273.
- Hossain, M. M., Sultana, A., and Purohit, N. (2020). Mental Health Outcomes of Quarantine and Isolation for Infection Prevention: A Systematic Umbrella Review of the Global Evidence. *SSRN Electronic Journal*.
- Hu, R. K., Zuo, Y., Ly, T., Wang, J., Meera, P., Wu, Y. E., and Hong, W. (2021). An amygdala-to-hypothalamus circuit for social reward. *Nature neuroscience*, 24(6):831–842.
- Hudlicka, E. (2003). To feel or not to feel: The role of affect in human-computer interaction. *International Journal of Human Computer Studies*, 59(1-2):1–32.
- Inderbitzin, M., Våljamäe, A., Calvo, J. M. B., Verschure, P. F., and Bernardet, U. (2011). Expression of emotional states during locomotion based on canonical parameters. In *2011 IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, pages 809–814. IEEE.
- Jacko, J. A. (2012). *Human computer interaction handbook: Fundamentals, evolving technologies, and emerging applications*. CRC press.

- Jacucci, G., Spagnolli, A., Freeman, J., and Gamberini, L. (2014). Symbiotic Interaction. *LNCS*, 8820:3–20.
- Jacucci, G., Spagnolli, A., Freeman, J., and Gamberini, L. (2015). Symbiotic interaction: a critical definition and comparison to other human-computer paradigms. In *International workshop on symbiotic interaction*, pages 3–20. Springer.
- James, W. (1884). What is an Emotion? *Mind*, 9:188–205.
- Jin, Y., Cardoso, B., and Verbert, K. (2017). How do different levels of user control affect cognitive load and acceptance of recommendations? In *CEUR Workshop Proceedings*, volume 1884, pages 35–42.
- John, O. P., Donahue, E. M., and Kentle, R. L. (1991). Big five inventory. *Journal of Personality and Social Psychology*.
- Johnson, K. T. and Picard, R. W. (2020). Advancing Neuroscience through Wearable Devices. *Neuron*, 108(1):8–12.
- Jones, F. (2006). Strategies to enhance chronic disease self-management: How can we apply this to stroke? *Disability and Rehabilitation*, 28(13-14):841–847.
- Jones, T. A. (2017). Motor compensation and its effects on neural reorganization after stroke. *Nature Reviews Neuroscience*, 18(5):267–280.
- Joyce, R. and Gupta, G. (1990). Identity authentication based on keystroke latencies. *Communications of the ACM*, 33(2):168–176.
- Ju, W. and Leifer, L. (2008). The design of implicit interactions: Making interactive systems less obnoxious. *Design Issues*, 24(3):72–84.
- Kawohl, W. and Nordt, C. (2020). Covid-19, unemployment, and suicide. *The Lancet Psychiatry*, 7(5):389–390.
- Keefer, K. V. (2015). Self-Report Assessments of Emotional Competencies: A Critical Look at Methods and Meanings. *Journal of Psychoeducational Assessment*.

- Kemp, A. H. and Quintana, D. S. (2013). The relationship between mental and physical health: insights from the study of heart rate variability. *International journal of Psychophysiology*, 89(3):288–296.
- Kerckhove, A. V. and Pandelaere, M. (2018). Why Are You Swiping Right? The Impact of Product Orientation on Swiping Responses. *Journal of Consumer Research*, 45(3):633–647.
- Kessels, R. P., Van Zandvoort, M. J., Postma, A., Kappelle, L. J., and De Haan, E. H. (2000). The Corsi Block-Tapping Task: Standardization and normative data. *Applied Neuropsychology*, 7(4):252–258.
- Kim, J., Thayabaranathan, T., Donnan, G. A., Howard, G., Howard, V. J., Rothwell, P. M., Feigin, V., Norrving, B., Owolabi, M., Pandian, J., Liu, L., Cadilhac, D. A., and Thrift, A. G. (2020). Global Stroke Statistics 2019. *International Journal of Stroke*, 15(8):819–838.
- Kim, J., Zhang, X., Muralidhar, S., LeBlanc, S. A., and Tonegawa, S. (2017). Basolateral to central amygdala neural circuits for appetitive behaviors. *Neuron*, 93(6):1464–1479.
- Kim, M.-K., Kim, M., Oh, E., and Kim, S.-P. (2013). A review on the Computational Methods for Emotional State Estimation from the Human EEG. *Computational and mathematical methods in medicine*, 2013.
- Kiss, M. and Eimer, M. (2008). ERPs reveal subliminal processing of fearful faces. *Psychophysiology*, 45(2):318–326.
- Klein, A. S., Dolensek, N., Weiand, C., and Gogolla, N. (2021). Fear balance is maintained by bodily feedback to the insular cortex in mice. *Science*, 374(6570):1010–1015.
- Ko, B. C. (2018). A brief review of facial emotion recognition based on visual information. *Sensors (Switzerland)*, 18(2).

- Kohlmorgen, J., Dornhege, G., Braun, M. L., Blankertz, B., Müller, K.-R., Curio, G., Hagemann, K., Bruns, A., Schrauf, M., and Kincses, W. (2007). Improving Human Performance in a Real Operating Environment through Real-Time Mental Workload Detection. In *Toward Brain-Computer Interfacing*, pages 409–422. MIT Press.
- Kolakowska, A. (2013). A review of emotion recognition methods based on keystroke dynamics and mouse movements. In *2013 6th International Conference on Human System Interactions (HSI)*., pages 548–555, Sopot, Poland.
- Kong, X., Wei, D., Li, W., Cun, L., Xue, S., Zhang, Q., and Qiu, J. (2014). Neuroticism and extraversion mediate the association between loneliness and the dorsolateral prefrontal cortex. *Experimental Brain Research*, 233(1):157–164.
- Kongthong, N., Minami, T., and Nakauchi, S. (2013). Semantic processing in subliminal face stimuli: An EEG and tDCS study. *Neuroscience Letters*, 544:141–146.
- Koolagudi, S. G. and Rao, K. S. (2012). Emotion recognition from speech: A review. *International Journal of Speech Technology*, 15(2):99–117.
- Kostoulas, T., Chanel, G., Muszynski, M., Lombardo, P., and Pun, T. (2017). Films, affective computing and aesthetic experience: Identifying emotional and aesthetic highlights from multimodal signals in a social setting. *Frontiers in ICT*, 4:11.
- Kothe, C. A. and Makeig, S. (2013). BCILAB: A platform for brain-computer interface development. *Journal of Neural Engineering*, 10(5).
- Kotsia, I., Zafeiriou, S., and Fotopoulos, S. (2013). Affective gaming: A comprehensive survey. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 663–670.

- Kouider, S. and Dehaene, S. (2007). Levels of processing during non-conscious perception: A critical review of visual masking. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481):857–875.
- Koutsiana, E., Ladakis, I., Fotopoulos, D., Chytas, A., Kilintzis, V., and Chouvarda, I. (2020). Serious gaming technology in upper extremity rehabilitation: Scoping review. *JMIR Serious Games*, 8(4).
- Kreibig, S. D. (2010). Autonomic nervous system activity in emotion: A review. *Biological psychology*, 84(3):394–421.
- Kroenke, K., Spitzer, R. L., and Williams, J. B. (2001). The PHQ-9: Validity of a brief depression severity measure. *Journal of General Internal Medicine*, 16(9):606–613.
- Krueger, A. B. and Mueller, A. I. (2012). Time use, emotional well-being, and unemployment: Evidence from longitudinal data. *American Economic Review*, 102(3):594–599.
- Krumhuber, E. G., Kappas, A., and Manstead, A. S. (2013). Effects of dynamic aspects of facial expressions: A review. *Emotion Review*, 5(1):41–46.
- Kuppens, P., Tuerlinckx, F., Russell, J. A., and Barrett, L. F. (2013). The relation between valence and arousal in subjective experience. *Psychological Bulletin*, 139(4):917–940.
- Kuppens, P., Tuerlinckx, F., Yik, M., Koval, P., Coosemans, J., Zeng, K. J., and Russell, J. A. (2017). The Relation Between Valence and Arousal in Subjective Experience Varies With Personality and Culture. *Journal of Personality*, 85(4):530–542.
- Kurdi, B., Lozano, S., and Banaji, M. R. (2017). Introducing the Open Affective Standardized Image Set (OASIS). *Behavior Research Methods*, 49(2):457–470.

- Lai, S. M., Studenski, S., Duncan, P. W., and Perera, S. (2002). Persisting consequences of stroke measured by the stroke impact scale. *Stroke*, 33(7):1840–1844.
- Lali, P., Naghizadeh, M., Nasrollahi, H., Moradi, H., and Mirian, M. S. (2014). Your mouse can tell about your emotions. In *Proceedings of the 4th International Conference on Computer and Knowledge Engineering, ICCKE 2014*, pages 47–51.
- Lanata, A., Armato, A., Valenza, G., and Scilingo, E. P. (2011). Eye tracking and pupil size variation as response to affective stimuli: A preliminary study. In *2011 5th International Conference on Pervasive Computing Technologies for Healthcare and Workshops, PervasiveHealth 2011*, pages 78–84.
- Lang, P. J., Bradley, M. M., and Cuthbert, B. N. (2008). International Affective Picture System (IAPS): Affective ratings of pictures and instruction manual. *Technical Report A-8. University of Florida, Gainesville, FL*.
- Lang, P. J., Greenwald, M. K., Bradley, M. M., and Hamm, A. O. (1993). Looking at pictures: Affective, facial, visceral, and behavioral reactions. *Psychophysiology*, 30(3):261–273.
- Lavie, N., Hirst, A., De Fockert, J. W., and Viding, E. (2004). Load theory of selective attention and cognitive control. *Journal of Experimental Psychology: General*, 133(3):339–354.
- Lazarus, R. S., Averill, J. R., and Opton Jr, E. M. (1970). Towards a cognitive theory of emotion. In *Feelings and emotions*, pages 207–232. Elsevier.
- Lazarus, R. S. and McCleary, R. A. (1951). Autonomic discrimination without awareness: a study of subception. *Psychological Review*, 58(2):113–122.

- LeDoux, J. E. and Brown, R. (2017). A higher-order theory of emotional consciousness. *Proceedings of the National Academy of Sciences of the United States of America*, 114(10):E2016–E2025.
- LeDoux, J. E. and Hofmann, S. G. (2018). The subjective experience of emotion: a fearful view. *Current Opinion in Behavioral Sciences*, 19:67–72.
- Lenaert, B., van Kampen, N., van Heugten, C., and Ponds, R. (2020). Real-time measurement of post-stroke fatigue in daily life and its relationship with the retrospective Fatigue Severity Scale. *Neuropsychological Rehabilitation*, 0(0):1–15.
- Lessiter, J., Freeman, J., Keogh, E., and Davidoff, J. (2001). A cross-media presence questionnaire: The ITC-sense of presence inventory. *Presence: Teleoperators and Virtual Environments*.
- Lessiter, J., Miotto, A., Freeman, J., Verschure, P., and Bernardet, U. (2011). CEEDs: Unleashing the Power of the Subconscious. *Procedia Computer Science*, 7:214–215.
- Levenson, R. W. (2014). The autonomic nervous system and emotion. *Emotion Review*, 6(2):100–112.
- Limanowski, J. and Blankenburg, F. (2016). Integration of Visual and Proprioceptive Limb Position Information in Human Posterior Parietal, Premotor, and Extrastriate Cortex. *Journal of Neuroscience*, 36(9):2582–2589.
- Lin, T., Hu, W., Imamiya, A., and Omata, M. (2006). Large display size enhances user experience in 3D games. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 4073 LNCS, pages 257–262.
- Lipka, J., Miltner, W. H. R., and Straube, T. (2011). Vigilance for threat interacts with amygdala responses to subliminal threat cues in specific phobia. *Biological Psychiatry*, 70(5):472–478.

- Lipton, Z. C., Elkan, C., and Naryanaswamy, B. (2014). Optimal thresholding of classifiers to maximize F1 measure. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 8725 LNAI, pages 225–239. Springer Verlag.
- Liu, D., Ren, Y., Yan, F., Li, Y., Xu, X., Yu, X., Qu, W., Wang, Z., Tian, B., Yang, F., Yao, Y., Tan, Y., Jiang, R., and Tan, S. (2020). Psychological Impact and Predisposing Factors of the Coronavirus Disease 2019 (COVID-19) Pandemic on General Public in China. *SSRN Electronic Journal*.
- Liu, Z., Xu, A., Guo, Y., Mahmud, J. U., Liu, H., and Akkiraju, R. (2018). Seemo: A computational approach to See emotions. *Conference on Human Factors in Computing Systems - Proceedings*, 2018-April.
- Lombard, M. and Ditton, T. (1997). At the Heart of It All: The Concept of Presence. *Journal of Computer-Mediated Communication*, 3(2).
- López-Carral, H., Blancas-Muñoz, M., Mura, A., Omedas, P., España-Cumellas, À., Martínez-Bueno, E., Milliken, N., Moore, P., Haque, L., Gilroy, S., and Verschure, P. F. M. J. (2022). A virtual reality system for the simulation of neurodiversity. In Yang, X.-S., Sherratt, S., Dey, N., and Joshi, A., editors, *Proceedings of Sixth International Congress on Information and Communication Technology*, pages 523–531, Singapore. Springer Singapore.
- López-Carral, H., de la Torre Costa, J., Freire, I. T., Mura, A., and Verschure, P. F. (2023a). A socially cooperative cognitive architecture for rehabilitation. *Manuscript in preparation*.
- López-Carral, H., Grechuta, K., and Verschure, P. F. (2020). Subjective ratings of emotive stimuli predict the impact of the COVID-19 quarantine on affective states. *PLoS ONE*, 15(8 August):1–15.

- López-Carral, H., Grechuta, K., and Verschure, P. F. (2023b). Implicit measures identify differential effects of COVID-19 on the psychological health of diverse populations. *Manuscript in preparation*.
- López-Carral, H., Omedas, P., and Verschure, P. F. (2023c). Experimental evaluation of an immersive, interactive and confluent system. *International Journal of Human-Computer Studies* (submitted).
- López-Carral, H., Omedas, P., Zucca, R., and Verschure, P. F. (2023d). A distributed sensing architecture for psychophysiological experimentation in virtual reality. *Manuscript in preparation*.
- López-Carral, H., Santos-Pata, D., Zucca, R., and Verschure, P. F. (2019). How you type is what you type: Keystroke dynamics correlate with affective content. *2019 8th International Conference on Affective Computing and Intelligent Interaction, ACII 2019*, pages 359–363.
- López-Carral, H. and Verschure, P. F. (2022). An Assistive System for Transferring Domain Knowledge to Novice Officers. *European Law Enforcement Research Bulletin*, 22(6).
- López-Carral, H. and Verschure, P. F. (2023). Binary affective ratings through swiping. *Manuscript in preparation*.
- López-Carral, H., Zucca, R., Omedas, P., and Verschure, P. F. (2023e). Induction of unconscious processing of emotional stimuli in an immersive system. *Manuscript in preparation*.
- Luxton, D. D., McCann, R. A., Bush, N. E., Mishkind, M. C., and Reger, G. M. (2011). MHealth for mental health: Integrating smartphone technology in behavioral healthcare. *Professional Psychology: Research and Practice*, 42(6):505–512.
- Magenat-Thalmann, N. and Zhang, Z. (2014). Assistive social robots for people with special needs. In *Proceedings of 2014 International Conference on Contemporary Computing and Informatics, IC3I 2014*, pages 1374–1380.

- Magnusson, C., Anastassova, M., Paneels, S., Rasmus-Grohn, K., Rydeman, B., Randall, G., Ortiz Fernandez, L., Bouilland, S., and Pager, J. (2018). Stroke and Universal Design. *Studies in health technology and informatics*, 256:854–861.
- Maier, M., Ballester, B. R., and Verschure, P. F. (2019). Principles of Neurorehabilitation After Stroke Based on Motor Learning and Brain Plasticity Mechanisms. *Frontiers in Systems Neuroscience*, 13(December):1–18.
- Maier, W., Philipp, M., and Gerken, A. (1985). Dimensions of the hamilton depression scale. factor analysis studies. *European archives of psychiatry and neurological sciences*, 234(6):417.
- Mäntylä, M. V., Graziotin, D., and Kuutila, M. (2018). The evolution of sentiment analysis-A review of research topics, venues, and top cited papers. *Computer Science Review*, 27:16–32.
- Mathews, Z., Bermudez i Badia, S., and Verschure, P. F. (2007). A novel brain-based approach for multi-modal multi-target tracking in a mixed reality space. *Intuition Workshop and Conference on Virtual Reality and Virtual Environments*, 3(1).
- Mathews, Z., Cetnarski, R., and Verschure, P. F. (2014). Visual anticipation biases conscious decision making but not bottom-up visual processing. *Frontiers in Psychology*, 5(OCT).
- Matthews, T., Danese, A., Wertz, J., Odgers, C. L., Ambler, A., Moffitt, T. E., and Arseneault, L. (2016). Social isolation, loneliness and depression in young adulthood: a behavioural genetic analysis. *Social Psychiatry and Psychiatric Epidemiology*, 51(3):339–348.
- Mauss, I. B. and Robinson, M. D. (2009). Measures of emotion: A review. *Cognition and emotion*, 23(2):209–237.
- Mayer, B. and Merckelbach, H. (1999). Unconscious processes, subliminal stimulation, and anxiety. *Clinical Psychology Review*, 19(5):571–590.

- Mazzocchi, F. (2015). Could Big Data be the end of theory in science? *EMBO reports*, 16(10):1250–1255.
- McCorry, L. K. (2007). Physiology of the autonomic nervous system. *American Journal of Pharmaceutical Education*, 71(4).
- McMahan, R. P., Gorton, D., Gresock, J., McConnell, W., and Bowman, D. A. (2006). Separating the effects of level of immersion and 3D interaction techniques. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology, VRST*.
- Mehrabian, A. (1980). *Basic dimensions for a general psychological theory: Implications for personality, social, environmental, and developmental studies*. Oelgeschlager, Gunn & Hain.
- Méndez-Bértolo, C., Moratti, S., Toledano, R., Lopez-Sosa, F., Martínez-Alvarez, R., Mah, Y. H., Vuilleumier, P., Gil-Nagel, A., and Strange, B. A. (2016). A fast pathway for fear in human amygdala. *Nature Neuroscience*, 19(8):1041–1049.
- Menghini, L., Gianfranchi, E., Cellini, N., Patron, E., Tagliabue, M., and Sarlo, M. (2019). Stressing the accuracy: Wrist-worn wearable sensor validation over different conditions. *Psychophysiology*, 56(11):1–15.
- Meyerbröker, K. and Emmelkamp, P. M. (2010). Virtual reality exposure therapy in anxiety disorders: A systematic review of process-and-outcome studies. *Depression and Anxiety*, 27(10):933–944.
- Miles, S. H. (2015). Kaci hickox: public health and the politics of fear. *The American Journal of Bioethics*, 15(4):17–19.
- Milgram, P., Takemura, H., Utsumi, A., and Kishino, F. (1994). Mixed Reality (MR) Reality-Virtuality (RV) Continuum. *Proceedings of SPIE - The International Society for Optical Engineering*, 2351(Telemanipulator and Telepresence Technologies):282–292.
- Milner, B. (1971). Interhemispheric differences in the localization of psychological processes in man. *British Medical Bulletin*, 27(3):272–277.

- Mitchell, P. and Ropar, D. (2004). Visuo-spatial abilities in autism: A review.
- Mohanavelu, K., Lamshe, R., Poonguzhali, S., Adalarasu, K., and Jagannath, M. (2017). Assessment of human fatigue during physical performance using physiological signals: A review. *Biomedical and Pharmacology Journal*, 10(4):1887–1896.
- Monrose, F. and Rubin, A. (1997). Authentication via keystroke dynamics. In *Proceedings of the 4th ACM conference on Computer and communications security*, pages 48–56. Association for Computing Machinery (ACM).
- Monrose, F. and Rubin, A. D. (2000). Keystroke dynamics as a biometric for authentication. *Future Generation Computer Systems*, 16(4):351–359.
- Moors, A., Ellsworth, P. C., Scherer, K. R., and Frijda, N. H. (2013). Appraisal theories of emotion: State of the art and future development. *Emotion Review*, 5(2):119–124.
- Moreno, R. (2005). Instructional technology: Promise and pitfalls. In *Technology-Based Education: Bringing Researchers and Practitioners Together*, pages 1–19. Information Age Publishing.
- Morgane, P. J., Galler, J. R., and Mokler, D. J. (2005). A review of systems and networks of the limbic forebrain/limbic midbrain.
- Moscarello, J. M. and LeDoux, J. E. (2013). The contribution of the amygdala to aversive and appetitive pavlovian processes. *Emotion Review*, 5(3):248–253.
- Muñoz-Saavedra, L., Miró-Amarante, L., and Domínguez-Morales, M. (2020). Augmented and virtual reality evolution and future tendency. *Applied sciences*, 10(1):322.

- Namburi, P., Beyeler, A., Yorozu, S., Calhoun, G. G., Halbert, S. A., Wichmann, R., Holden, S. S., Mertens, K. L., Anahtar, M., Felix-Ortiz, A. C., et al. (2015). A circuit mechanism for differentiating positive and negative associations. *Nature*, 520(7549):675–678.
- Nandwani, P. and Verma, R. (2021). A review on sentiment analysis and emotion detection from text. *Social Network Analysis and Mining*, 11(1):1–19.
- Narayan, M., Waugh, L., Zhang, X., Bafna, P., and Bowman, D. (2006). Quantifying the benefits of immersion for collaboration in Virtual Environments. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology, VRST*, volume 2006, pages 78–81.
- Ni, T., Bowman, D. A., and Chen, J. (2006). Increased display size and resolution improve task performance in information-rich virtual environments. In *Proceedings - Graphics Interface*, volume 2006, pages 139–146.
- Nirme, J., Duff, A., and Verschure, P. F. (2011). Adaptive rehabilitation gaming system: On-line individualization of stroke rehabilitation. *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, pages 6749–6752.
- Niven, K. (2013). Affect. In Gellman, Marc D. and Turner, J. R., editor, *Encyclopedia of Behavioral Medicine*, pages 49–50. Springer New York, New York, NY.
- Nobles, J., Martin, F., Dawson, S., Moran, P., and Savovic, J. (2020). The potential impact of COVID-19 on mental health outcomes and the implications for service solutions. *Bristol: National Institute for Health Research, University of Bristol*.
- Noroozi, F., Corneanu, C. A., Kaminska, D., Sapinski, T., Escalera, S., and Anbarjafari, G. (2021). Survey on Emotional Body Gesture Recognition. *IEEE Transactions on Affective Computing*, 12(2):505–523.

- Ojeda, A., Bigdely-Shamlo, N., and Makeig, S. (2014). MoBILAB: an open source toolbox for analysis and visualization of mobile brain/body imaging data. *Frontiers in Human Neuroscience*, 8.
- Omedas, P., Betella, A., Zucca, R., Arsiwalla, X. D., Pacheco, D., Wagner, J., Lingenfelter, F., Andre, E., Mazzei, D., Lanatá, A., Tognetti, A., de Rossi, D., Grau, A., Goldhoorn, A., Guerra, E., Alquezar, R., Sanfeliu, A., and Verschure, P. F. M. J. (2014). XIM-engine: A Software Framework to Support the Development of Interactive Applications That Uses Conscious and Unconscious Reactions in Immersive Mixed Reality. *Proceedings of the 2014 Virtual Reality International Conference*, pages 26:1—26:4.
- Paradiso, R., Loriga, G., and Taccini, N. (2005). A wearable health care system based on knitted integrated sensors. *IEEE Transactions on Information Technology in Biomedicine*, 9(3):337–344.
- Park, D. H., Kim, H. K., Choi, I. Y., and Kim, J. K. (2012). A literature review and classification of recommender systems research. *Expert Systems with Applications*, 39(11):10059–10072.
- Park, M., Naaman, M., and Berger, J. (2016). A data-driven study of view duration on YouTube. In *Proceedings of the 10th International Conference on Web and Social Media, ICWSM 2016*, pages 651–654.
- Parsons, T. D. (2015). Virtual Reality for Enhanced Ecological Validity and Experimental Control in the Clinical, Affective and Social Neurosciences. *Frontiers in Human Neuroscience*, 9(December):1–19.
- Patrick, E., Cosgrove, D., Slavkovic, A., Rode, J. A., Verratti, T., and Chiselko, G. (2000). Using a large projection screen as an alternative to head-mounted displays for virtual environments. In *Conference on Human Factors in Computing Systems - Proceedings*, pages 478–485.
- Paulhus, D. L. (2002). Socially desirable responding: The evolution of a construct. *The role of constructs in psychological and educational measurement*, 49459.

- Paulhus, D. L. (2017). Socially desirable responding on self-reports. *Encyclopedia of personality and individual differences*, pages 1–5.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Philippot, P., Chapelle, G., and Blairy, S. (2002). Respiratory feedback in the generation of emotion. *Cognition and Emotion*, 16(5):605–627.
- Picard, R. W. (1995). Affective Computing. Technical report, MIT Media Laboratory.
- Picard, R. W. (1997). *Affective Computing*. The MIT Press.
- Picard, R. W., Vyzas, E., and Healey, J. (2001). Toward machine emotional intelligence: analysis of affective physiological state. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(10):1175–1191.
- Pinto, J. F., Carvalho, H. R., Chambel, G. R., Ramiro, J., and Gonçalves, A. (2018). Adaptive gameplay and difficulty adjustment in a gamified upper-limb rehabilitation. *2018 IEEE 6th International Conference on Serious Games and Applications for Health, SeGAH 2018*, pages 1–8.
- Polys, N. F., Kim, S., and Bowman, D. A. (2007). Effects of information layout, screen size, and field of view on user performance in information-rich virtual environments. *Computer Animation and Virtual Worlds*, 18(1):19–38.
- Pomplun, M. and Sunkara, S. (2003). Pupil Dilation as an Indicator of Cognitive Workload in Human-Computer Interaction. *Human-centered computing: Cognitive, social and ergonomic aspects. Vol. 3 of the Proceedings of the 10th International Conference on Human-Computer Interaction*, pages 542–546.

- Porges, S. W. et al. (1997). Emotion: An evolutionary by-product of the neural regulation of the autonomic nervous system. *The integrative neurobiology of affiliation*, 807:62–67.
- Posner, J., Russell, J. A., and Peterson, B. S. (2005). The circumplex model of affect: An integrative approach to affective neuroscience, cognitive development, and psychopathology. *Dev Psychopathol*, 17(3):715–734.
- Ragan, E. D., Sowndararajan, A., Kopper, R., and Bowman, D. (2010). The effects of higher levels of immersion on procedure memorization performance and implications for educational virtual environments. *Presence: Teleoperators and Virtual Environments*, 19(6):527–543.
- Rajkumar, R. P. (2020). COVID-19 and mental health: A review of the existing literature. *Asian Journal of Psychiatry*, 52.
- Rajmohan, V. and Mohandas, E. (2007). The limbic system. *Indian Journal of Psychiatry*, 49(2):132.
- Rammstedt, B. and John, O. P. (2007). Measuring personality in one minute or less: A 10-item short version of the Big Five Inventory in English and German. *Journal of Research in Personality*, 41(1):203–212.
- Reason, J. T. and Brand, J. J. (1975). *Motion Sickness*. Academic Press.
- Reinertsen, E. and Clifford, G. D. (2018). A review of physiological and behavioral monitoring with digital sensors for neuropsychiatric illnesses. *Physiological measurement*, 39(5):05TR01.
- Ren, X., Silpasuwanchai, C., and Cahill, J. (2019). Human-engaged computing: the future of human–computer interaction. *CCF transactions on pervasive computing and interaction*, 1(1):47–68.
- Renard, Y., Lotte, F., Gibert, G., Congedo, M., Maby, E., Delannoy, V., Bertrand, O., and Lécuyer, A. (2010). OpenViBE: An open-source

- software platform to design, test, and use brain-computer interfaces in real and virtual environments. *Presence: Teleoperators and Virtual Environments*, 19(1):35–53.
- Ricci, F., Rokach, L., Shapira, B., Kantor, P. B., and Ricci, F. (2011). *Recommender Systems Handbook*. Springer New York, NY.
- Richthammer, C. and Pernul, G. (2020). Situation awareness for recommender systems. *Electronic Commerce Research*, 20(4):783–806.
- Riess, H. (2017). The Science of Empathy. *Journal of Patient Experience*, 4(2):74–77.
- Rinaldi, C., Attanasio, M., Valenti, M., Mazza, M., and Keller, R. (2021). Autism spectrum disorder and personality disorders: Comorbidity and differential diagnosis. *World Journal of Psychiatry*, 11(12):1366.
- Robinson, R., Wiley, K., Rezaeivahdati, A., Klarkowski, M., and Mandryk, R. L. (2020). ” let’s get physiological, physiological!” a systematic review of affective gaming. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*, pages 132–147.
- Roca, M., Gili, M., Garcia-Campayo, J., and García-Toro, M. (2013). Economic crisis and mental health in Spain. *The Lancet*, 382(9909):1977–1978.
- Rothstein, M. A., Alcalde, G., Elster, N. R., Majumder, M. A., Palmer, L. I., Stone, H. T., and Hoffman, R. E. (2003). *Quarantine and isolation: Lessons learned from SARS*. Diane Pub Co.
- Rottenberg, J., Gross, J. J., and Gotlib, I. H. (2005). Emotion context insensitivity in major depressive disorder. *Journal of abnormal psychology*, 114(4):627.
- Rubin, G. J. and Wessely, S. (2020). The psychological effects of quarantining a city. *Bmj*, 368.

- Russell, J. A. (1979). Affective space is bipolar. *Journal of Personality and Social Psychology*, 37(3):345–356.
- Russell, J. A. (1980). A circumplex model of affect. *Journal of Personality and Social Psychology*, 39(6):1161–1178.
- Russell, J. A. (2003). Core Affect and the Psychological Construction of Emotion. *Psychological Review*, 110(1):145–172.
- Ryan, R. M. (1982). Control and information in the intrapersonal sphere: An extension of cognitive evaluation theory. *Journal of Personality and Social Psychology*, 43(3):450–461.
- Santini, Z. I., Jose, P. E., York Cornwell, E., Koyanagi, A., Nielsen, L., Hinrichsen, C., Meilstrup, C., Madsen, K. R., and Koushede, V. (2020). Social disconnectedness, perceived isolation, and symptoms of depression and anxiety among older Americans (NSHAP): a longitudinal mediation analysis. *The Lancet Public Health*, 5(1):e62–e70.
- Santos, O. C. (2016). Emotions and Personality in Adaptive e-Learning Systems: An Affective Computing Perspective. In Tkalčič M., De Carolis B., de Gemmis M., Odić A., K. A., editor, *Emotions and Personality in Personalized Services. Human-Computer Interaction Series*, pages 263–285. Springer.
- Santos-Pata, D., Zucca, R., López-Carral, H., and Verschure, P. F. (2019). Modulating grid cell scale and intrinsic frequencies via slow high-threshold conductances: A simplified model. *Neural Networks*.
- Sayers, J. (2001). The world health report 2001-mental health: new understanding, new hope. *Bulletin of the World Health Organization*, 79:1085–1085.
- Schaaff, K., Degen, R., Adler, N., and Adam, M. T. (2012). Measuring affect using a standard mouse device. *Biomedizinische Technik*, 57(SUPPL. 1 TRACK-N):761–764.

- Schalk, G., McFarland, D. J., Hinterberger, T., Birbaumer, N., and Wolpaw, J. R. (2004). BCI2000: A general-purpose brain-computer interface (BCI) system. *IEEE Transactions on Biomedical Engineering*, 51(6):1034–1043.
- Scharlau, I., Ansorge, U., and Breitmeyer, B. G. (2006). Trends and styles in visual masking. *Advances in Cognitive Psychology*, 2(1):1–5.
- Scherer, K. R. (1999). Appraisal theory. In *Handbook of cognition and emotion*, pages 637–663. John Wiley & Sons Ltd, Hoboken, NJ, US.
- Schmidt, F., Haberkamp, A., and Schmidt, T. (2011). Dos and don'ts in response priming research. *Advances in Cognitive Psychology*, 7(1):120–131.
- Schmitter, A. M. (2016). 17th and 18th Century Theories of Emotions. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 201 edition.
- Schuchardt, P. and Bowman, D. A. (2007). The benefits of immersion for spatial understanding of complex underground cave systems. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology, VRST*, pages 121–124.
- Schuller, D. M. and Schuller, B. W. (2021). A Review on Five Recent and Near-Future Developments in Computational Processing of Emotion in the Human Voice. *Emotion Review*, 13(1):44–50.
- Segerstrom, S. C. and Nes, L. S. (2007). Heart rate variability reflects self-regulatory strength, effort, and fatigue. *Psychological Science*, 18(3).
- Selvaraj, J., Murugappan, M., Wan, K., and Yaacob, S. (2013). Classification of emotional states from electrocardiogram signals: A non-linear approach based on hurst. *BioMedical Engineering Online*, 12(1).
- Serim, B. and Jacucci, G. (2019). Explicating “implicit interaction”: An examination of the concept and challenges for research. *Conference on Human Factors in Computing Systems - Proceedings*, pages 1–16.

- Serrano-Guerrero, J., Olivas, J. A., Romero, F. P., and Herrera-Viedma, E. (2015). Sentiment analysis: A review and comparative analysis of web services. *Information Sciences*, 311:18–38.
- Setz, C., Arnrich, B., Schumm, J., La Marca, R., Tröster, G., and Ehlert, U. (2009). Discriminating stress from cognitive load using a wearable eda device. *IEEE Transactions on information technology in biomedicine*, 14(2):410–417.
- Shah, J. N., Shah, J., and Shah, J. N. (2020). Quarantine, isolation and lockdown: in context of COVID-19. *Journal of Patan Academy of Health Sciences*, 7(1):48–57.
- Shamay-Tsoory, S. G., Aharon-Peretz, J., and Perry, D. (2009). Two systems for empathy: A double dissociation between emotional and cognitive empathy in inferior frontal gyrus versus ventromedial prefrontal lesions. *Brain*, 132(3):617–627.
- Shankar, A., McMunn, A., Banks, J., and Steptoe, A. (2011). Loneliness, Social Isolation, and Behavioral and Biological Health Indicators in Older Adults. *Health Psychology*, 30(4):377–385.
- Sharpley, S., Cobb, S., Moody, A., and Wilson, J. R. (2008). Virtual reality induced symptoms and effects (VRISE): Comparison of head mounted display (HMD), desktop and projection display systems. *Displays*, 29(2):58–69.
- Shim, K. C., Park, J. S., Kim, H. S., Kim, J. H., Park, Y. C., and Ryu, H. I. (2003). Application of virtual reality technology in biology education. *Journal of Biological Education*, 37(2):71–74.
- Shu, L., Xie, J., Yang, M., Li, Z., Li, Z., Liao, D., Xu, X., and Yang, X. (2018). A review of emotion recognition using physiological signals. *Sensors*, 18(7).
- Simons, R. F., Detenber, B. H., Reiss, J. E., and Shults, C. W. (2000). Image motion and context: A between- and within-subjects comparison. *Psychophysiology*, 37(5):706–710.

- Slater, M., Linakis, V., Usoh, M., Kooper, R., and Street, G. (1996). Immersion, Presence, and Performance in Virtual Environments: An Experiment with Tri-Dimensional Chess. In *ACM Virtual Reality Software and Technology (VRST)*, pages 163–172.
- Sloan, D. M., Bradley, M. M., Dimoulas, E., and Lang, P. J. (2002). Looking at facial expressions: Dysphoria and facial emg. *Biological psychology*, 60(2-3):79–90.
- Sloan, D. M., Strauss, M. E., Quirk, S. W., and Sajatovic, M. (1997). Subjective and expressive emotional responses in depression. *Journal of affective disorders*, 46(2):135–141.
- Smarr, K. L. and Keefer, A. L. (2011). Measures of depression and depressive symptoms: beck depression inventory-II (BDI-II), Center for Epidemiologic Studies Depression Scale (CES-D), geriatric depression scale (GDS), hospital anxiety and depression scale (HADS), and patient health Questionnaire-9 (PHQ-9). *Arthritis care & research*, 63(S11):S454–S466.
- Soleymani, M., Asghari-Esfeden, S., Fu, Y., and Pantic, M. (2016). Analysis of EEG Signals and Facial Expressions for Continuous Emotion Detection. *IEEE Transactions on Affective Computing*.
- Song, Y., Dixon, S., and Pearce, M. T. (2012). A survey of music recommendation systems and future perspectives. In *Proceedings of the 9th International Symposium on Computer Music Modeling and Retrieval (CMMR)*, volume 4, pages 395–410.
- Sowndararajan, A., Wang, R., and Bowman, D. A. (2008). Quantifying the benefits of immersion for procedural training. In *IPT/EDT 2008 - Immersive Projection Technologies/Emerging Display Technologies Workshop*.
- Sprang, G. and Silman, M. (2013). Posttraumatic stress disorder in parents and youth after health-related disasters. *Disaster medicine and public health preparedness*, 7(1):105–110.

- Sreeja, P. and Mahalakshmi, G. (2017). Emotion Models: A Review. *International Journal of Control Theory and Applications*, 10(8):651–657.
- Stanners, R. F., Coulter, M., Sweet, A. W., and Murphy, P. (1979). The pupillary response as an indicator of arousal and cognition. *Motivation and Emotion*, 3(4):319–340.
- Steultjens, E. M., Dekker, J., Bouter, L. M., Van de Nes, J. C., Cup, E. H., and Van den Ende, C. H. (2003). Occupational therapy for stroke patients: A systematic review. *Stroke*, 34(3):676–686.
- Stevens, E., Emmett, E., Wang, Y., McKeivitt, C., and Wolfe, C. (2017). *The burden of stroke in Europe*, volume 53. Stroke Alliance for Europe.
- Stocks, E. and Lishner, D. (2012). Empathy. In Ramachandran, V., editor, *Encyclopedia of Human Behavior (Second Edition)*, pages 32–37. Academic Press, San Diego, second edition edition.
- Stone, A. A., Bachrach, C. A., Jobe, J. B., Kurtzman, H. S., and Cain, V. S. (1999). *The science of self-report: Implications for research and practice*. Psychology Press.
- Stuber, G. D., Sparta, D. R., Stamatakis, A. M., van Leeuwen, W. A., Hardjoprajitno, J. E., Cho, S., Tye, K. M., Kempadoo, K. A., Zhang, F., Deisseroth, K., et al. (2011). Excitatory transmission from the amygdala to nucleus accumbens facilitates reward seeking. *Nature*, 475(7356):377–380.
- Sust, P. P., Solans, O., Fajardo, J. C., Peralta, M. M., Rodenas, P., Gabaldà, J., Eroles, L. G., Comella, A., Muñoz, C. V., Ribes, J. S., et al. (2020). Turning the crisis into an opportunity: digital health strategies deployed during the covid-19 outbreak. *JMIR Public Health and Surveillance*, 6(2):e19106.
- Sylaiou, S., Mania, K., Karoulis, A., and White, M. (2010). Exploring the relationship between presence and enjoyment in a virtual museum. *International Journal of Human Computer Studies*, 68(5):243–253.

- Szwoch, W. (2015). Emotion Recognition Using Physiological Signals. In *Proceedings of the Multimedia, Interaction, Design and Innovation - MIDI '15*, pages 1–8.
- Tamietto, M. and de Gelder, B. (2010). Neural bases of the non-conscious perception of emotional signals. *Nature Reviews Neuroscience*, 11(10):697–709.
- Taub, E., Crago, J. E., Burgio, L. D., Groomes, T. E., Cook, E. W., DeLuca, S. C., and Miller, N. E. (1994). an Operant Approach To Rehabilitation Medicine: Overcoming Learned Nonuse By Shaping. *Journal of the Experimental Analysis of Behavior*, 61(2):281–293.
- Taylor, D. A., Altman, I., Wheeler, L., and Kushner, E. N. (1969). Personality factors related to response to social isolation and confinement. *Journal of Consulting and Clinical Psychology*, 33(4):411–419.
- Teh, P. S., Teoh, A. B. J., and Yue, S. (2013). A Survey of Keystroke Dynamics Biometrics. *The Scientific World Journal*, 2013:1–24.
- Terranova, J. I., Yokose, J., Osanai, H., Marks, W. D., Yamamoto, J., Ogawa, S. K., and Kitamura, T. (2022). Hippocampal-amygdala memory circuits govern experience-dependent observational fear. *Neuron*, 110(8):1416–1431.
- Thayer, J. F. and Lane, R. D. (2009). Claude Bernard and the heart-brain connection: Further elaboration of a model of neurovisceral integration. *Neuroscience and Biobehavioral Reviews*, 33(2).
- The Center for Systems Science and Engineering, Johns Hopkins (2020). Coronavirus COVID-19 global cases. <https://www.arcgis.com/apps/opsdashboard/index.html#/bda7594740fd40299423467b48e9ecf6>.
- Timmermans, A. A., Seelen, H. A., Willmann, R. D., and Kingma, H. (2009). Technology-assisted training of arm-hand skills in stroke: Concepts on reacquisition of motor control and therapist guidelines for re-

- habilitation technology design. *Journal of NeuroEngineering and Rehabilitation*, 6(1).
- Tognetti, A., Bartalesi, R., Lorussi, F., and de Rossi, D. (2007). Body segment position reconstruction and posture classification by smart textiles. *Transactions of the Institute of Measurement & Control*, 29(3-4):215–253.
- Torales, J., O’Higgins, M., Castaldelli-Maia, J. M., and Ventriglio, A. (2020). The outbreak of COVID-19 coronavirus and its impact on global mental health. *International Journal of Social Psychiatry*, 66(4):317–320.
- Torres, E. P., Torres, E. A., Hernández-Álvarez, M., and Yoo, S. G. (2020). EEG-based BCI emotion recognition: A survey. *Sensors (Switzerland)*, 20(18):1–36.
- Tsai, C. W., Lai, C. F., Chao, H. C., and Vasilakos, A. V. (2015). Big data analytics: a survey. *Journal of Big Data*, 2(1).
- Tyng, C. M., Amin, H. U., Saad, M. N., and Malik, A. S. (2017). The influences of emotion on learning and memory. *Frontiers in Psychology*, 8(AUG).
- Umberson, D. and Karas Montez, J. (2010). Social Relationships and Health: A Flashpoint for Health Policy. *Journal of Health and Social Behavior*, 51(1_suppl):S54–S66.
- Vallila-Rohter, S. and Kiran, S. (2013). Non-linguistic learning and aphasia: Evidence from a paired associate and feedback-based task. *Neuropsychologia*, 51(1):79–90.
- Valtorta, N. K., Kanaan, M., Gilbody, S., Ronzi, S., and Hanratty, B. (2016). Loneliness and social isolation as risk factors for coronary heart disease and stroke: Systematic review and meta-analysis of longitudinal observational studies. *Heart*, 102(13):1009–1016.

- Van De Port, I. G., Kwakkel, G., Van Wijk, I., and Lindeman, E. (2006). Susceptibility to deterioration of mobility long-term after stroke: A prospective cohort study. *Stroke*, 37(1):167–171.
- Verbert, K., Manouselis, N., Ochoa, X., Wolpers, M., Drachler, H., Bosnic, I., and Duval, E. (2012). Context-aware recommender systems for learning: a survey and future challenges. *IEEE Transactions on Learning Technologies*, 5(4):318–335.
- Verschure, P. F. (2012). Distributed Adaptive Control: A theory of the Mind, Brain, Body Nexus. *Biologically Inspired Cognitive Architectures*, 1:55–72.
- Verschure, P. F., Pennartz, C. M., and Pezzulo, G. (2014). The why, what, where, when and how of goal-directed choice: neuronal and computational principles. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1655):20130483.
- Verschure, P. F., Voegtlin, T., and Douglas, R. J. (2003). Environmentally mediated synergy between perception and behaviour in mobile robots. *Nature*, 425(6958):620–624.
- Verschure, P. F. M. J. (2016). Synthetic consciousness: the distributed adaptive control perspective. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1701):20150448.
- Vigo, D., Thornicroft, G., and Atun, R. (2016). Estimating the true global burden of mental illness. *The Lancet Psychiatry*, 3(2):171–178.
- Villegas, N. M., Sánchez, C., Díaz-Cely, J., and Tamura, G. (2018). Characterizing context-aware recommender systems: A systematic literature review. *Knowledge-Based Systems*, 140:173–200.
- Vorberg, D., Mattler, U., Heinecke, A., Schmidt, T., and Schwarzbach, J. (2003). Different time courses for visual perception and action priming. *Proceedings of the National Academy of Sciences*, 100(10):6275–6280.

- Vouloutsis, V. and Verschure, P. F. (2018). Emotions and self-regulation. In Prescott, T. J., Lepora, N., and Verschure, P. F., editors, *Living Machines: A handbook of research in biomimetics and biohybrid systems*, chapter 34. Oxford University Press.
- Wafa, H. A., Wolfe, C. D., Emmett, E., Roth, G. A., Johnson, C. O., and Wang, Y. (2020). Burden of Stroke in Europe: Thirty-Year Projections of Incidence, Prevalence, Deaths, and Disability-Adjusted Life Years. *Stroke*, 51(8):2418–2427.
- Wagner, J., Lingenfelter, F., and André, E. (2011). The Social Signal Interpretation Framework (SSI) for Real Time Signal Processing and Recognition. In *Proceedings of Interspeech 2011*.
- Wagner, J., Lingenfelter, F., André, E., Mazzei, D., Tognetti, A., Lanatà, A., Rossi, D. D., Betella, A., Zucca, R., Omedas, P., and Verschure, P. F. M. J. (2013). A sensing architecture for empathetic data systems. In *Proceedings of the 4th Augmented Human International Conference, AH '13*, pages 96–99, New York, NY, USA. ACM.
- Wahis, J., Baudon, A., Althammer, F., Kerspern, D., Goyon, S., Hagiwara, D., Lefevre, A., Barteczko, L., Boury-Jamot, B., Bellanger, B., et al. (2021). Astrocytes mediate the effect of oxytocin in the central amygdala on neuronal activity and affective states in rodents. *Nature neuroscience*, 24(4):529–541.
- Walker, B. N. (2002). Magnitude estimation of conceptual data dimensions for use in sonification. *Journal of Experimental Psychology: Applied*, 8(4):211–221.
- Waltemate, T., Gall, D., Roth, D., Botsch, M., and Latoschik, M. E. (2018). The impact of avatar personalization and immersion on virtual body ownership, presence, and emotional response. *IEEE Transactions on Visualization and Computer Graphics*, 24(4):1643–1652.
- Wang, C., Pan, R., Wan, X., Tan, Y., Xu, L., Ho, C. S., and Ho, R. C. (2020a). Immediate psychological responses and associated factors

- during the initial stage of the 2019 coronavirus disease (COVID-19) epidemic among the general population in China. *International Journal of Environmental Research and Public Health*, 17(5).
- Wang, C.-A., Baird, T., Huang, J., Coutinho, J. D., Brien, D. C., and Munoz, D. P. (2018). Arousal Effects on Pupil Size, Heart Rate, and Skin Conductance in an Emotional Face Task. *Frontiers in Neurology*, 9.
- Wang, C.-M. and Chen, Y.-C. (2020). Design of an interactive mind calligraphy system by affective computing and visualization techniques for real-time reflections of the writer's emotions. *Sensors*, 20(20):5741.
- Wang, Y., Di, Y., Ye, J., and Wei, W. (2020b). Study on the public psychological states and its related factors during the outbreak of coronavirus disease 2019 (COVID-19) in some regions of China. *Psychology, Health and Medicine*.
- Wang, Y., Song, W., Tao, W., Liotta, A., Yang, D., Li, X., Gao, S., Sun, Y., Ge, W., Zhang, W., and Zhang, W. (2022). A systematic review on affective computing: emotion models, databases, and recent advances. *Information Fusion*, 83-84(March):19–52.
- Watson, D., Clark, L. A., and Tellegen, A. (1988). Development and Validation of Brief Measures of Positive and Negative Affect: The PANAS Scales. *Journal of Personality and Social Psychology*, 54(6):1063–1070.
- Whelan, R. (2008). Effective analysis of reaction time data. *Psychological Record*, 58(3):475–482.
- Wilson, J. M., Lee, J., and Shook, N. J. (2021). Covid-19 worries and mental health: the moderating effect of age. *Aging & Mental Health*, 25(7):1289–1296.
- Wolf, S. L., Lecraw, D. E., Barton, L. A., and Jann, B. B. (1989). Forced use of hemiplegic upper extremities to reverse the effect of learned

- nonuse among chronic stroke and head-injured patients. *Experimental Neurology*, 104(2):125–132.
- Woodward, K., Kanjo, E., Brown, D., McGinnity, T. M., Inkster, B., MacIntyre, D., and Tsanas, T. (2020). Beyond Mobile Apps: a Survey of Technologies for Mental Well-being. *IEEE Transactions on Affective Computing*.
- World Health Organization (2020). Novel coronavirus - China. <https://www.who.int/csr/don/12-january-2020-novel-coronavirus-china/en/>.
- Wu, S., Rizoiu, M. A., and Xie, L. (2018). Beyond views: Measuring and predicting engagement in online videos. In *12th International AAAI Conference on Web and Social Media, ICWSM 2018*, pages 434–443.
- Wundt, W. (1897). *Outlines of psychology*. W. Engelmann.
- Xiu, J., Zhang, Q., Zhou, T., Zhou, T.-t., Chen, Y., and Hu, H. (2014). Visualizing an emotional valence map in the limbic forebrain by tai-fish. *Nature neuroscience*, 17(11):1552–1559.
- Yadegaridehkordi, E., Noor, N. F. B. M., Ayub, M. N. B., Affal, H. B., and Hussin, N. B. (2019). Affective computing in education: A systematic review and future research. *Computers & Education*, 142:103649.
- Yang, L. and Qin, S.-f. (2021). A Review of Emotion Recognition Methods from Keystroke, Mouse, and Touchscreen Dynamics. *IEEE Access*, PP:1.
- Yannakakis, G. N. (2018). Enhancing health care via affective computing. *Malta Journal of Health Sciences*.
- Yazdani, A., Lee, J. S., and Ebrahimi, T. (2009). Implicit emotional tagging of multimedia using EEG signals and brain computer interface. In *1st ACM SIGMM International Workshop on Social Media, WSM'09, Co-located with the 2009 ACM International Conference on Multimedia, MM'09*, pages 81–88.

- Zander, T. O., Kothe, C., Jatzev, S., and Gaertner, M. (2010). Enhancing Human-Computer Interaction with Input from Active and Passive Brain-Computer Interfaces. In *Brain-Computer Interfaces: Applying our Minds to Human-Computer Interaction*, pages 181–199. Springer.
- Zelenski, J. M., Sobocko, K., and Whelan, D. C. (2013). Introversion, Solitude, and Subjective Well-Being. In Coplan, R. J. and Bowker, J. C., editors, *The Handbook of Solitude*, pages 184–201. John Wiley & Sons, Ltd.
- Zepf, S., Hernandez, J., Schmitt, A., Minker, W., and Picard, R. W. (2020). Driver emotion recognition for intelligent vehicles: A survey. *ACM Computing Surveys (CSUR)*, 53(3):1–30.
- Zhang, X., Chen, X., Li, Y., Lantz, V., Wang, K., and Yang, J. (2011). A framework for hand gesture recognition based on accelerometer and emg sensors. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 41(6):1064–1076.
- Zhao, J., Li, H., Kung, D., Fisher, M., Shen, Y., and Liu, R. (2020). Impact of the COVID-19 Epidemic on Stroke Care and Potential Solutions. *Stroke*, 51(7):1996–2001.
- Zheng, W. L., Guo, H. T., and Lu, B. L. (2015). Revealing critical channels and frequency bands for emotion recognition from EEG with deep belief network. *International IEEE/EMBS Conference on Neural Engineering, NER*, 2015-July:154–157.
- Zhu, J., Ji, L., and Liu, C. (2019). Heart rate variability monitoring for emotion and disorders of emotion. *Physiological measurement*, 40(6):064004.
- Zhu, N., Zhang, D., Wang, W., Li, X., Yang, B., Song, J., Zhao, X., Huang, B., Shi, W., Lu, R., et al. (2020). A novel coronavirus from patients with pneumonia in china, 2019. *New England Journal of Medicine*.

Zimmerman, M., Posternak, M. A., and Chelminski, I. (2005). Is it time to replace the hamilton depression rating scale as the primary outcome measure in treatment studies of depression? *Journal of Clinical Psychopharmacology*, 25(2):105–110.