

User-Centered Collaborative Visualization

Linnaeus University Dissertations

No 224/2015

USER-CENTERED COLLABORATIVE
VISUALIZATION

DANIEL CERNEA

LINNAEUS UNIVERSITY PRESS

User-Centered Collaborative Visualization

Doctoral dissertation, Department of Computer Science, Linnaeus University,
Växjö, Sweden, 2015

ISBN: 978-91-87925-65-8

Published by: Linnaeus University Press, 351 95 Växjö

Printed by: Elanders Sverige AB, 2015

This research has been performed during my doctoral studies at the University of Kaiserslautern, Germany, and at the Linnaeus University, Sweden. All aspects of this thesis were executed under the supervision of Prof. Dr. Achim Ebert and Prof. Dr. Andreas Kerren, and in close collaboration with their two research groups. This doctoral thesis is in form and content identical at both universities.

The thesis is also available as permanent link on the publication server of the University of Kaiserslautern:
<http://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:hbz:386-kluedo-40511>

Abstract

Cernea, Daniel (2015). *User-Centered Collaborative Visualization*, Linnaeus University Dissertation No 224/2015, ISBN: 978-91-87925-65-8. Written in English.

The last couple of years have marked the entire field of information technology with the introduction of a new global resource, called data. Certainly, one can argue that large amounts of information and highly interconnected and complex datasets were available since the dawn of the computer and even centuries before. However, it has been only a few years since digital data has exponentially expanded, diversified and interconnected into an overwhelming range of domains, generating an entire universe of zeros and ones. This universe represents a source of information with the potential of advancing a multitude of fields and sparking valuable insights. In order to obtain this information, this data needs to be explored, analyzed and interpreted.

While a large set of problems can be addressed through automatic techniques from fields like artificial intelligence, machine learning or computer vision, there are various datasets and domains that still rely on the human intuition and experience in order to parse and discover hidden information. In such instances, the data is usually structured and represented in the form of an interactive visual representation that allows users to efficiently explore the data space and reach valuable insights. However, the experience, knowledge and intuition of a single person also has its limits. To address this, collaborative visualizations allow multiple users to communicate, interact and explore a visual representation by building on the different views and knowledge blocks contributed by each person.

In this dissertation, we explore the potential of subjective measurements and user emotional awareness in collaborative scenarios as well as support flexible and user-centered collaboration in information visualization systems running on tabletop displays. We commence by introducing the concept of user-centered collaborative visualization (UCCV) and highlighting the context in which it applies. We continue with a thorough overview of the state-of-the-art in the areas of collaborative information visualization, subjectivity measurement and emotion visualization, combinable tabletop tangibles, as well as browsing history visualizations. Based on a new web browser history visualization for exploring user parallel browsing behavior, we introduce two novel user-centered techniques for supporting collaboration in co-located visualization systems. To begin with, we inspect the particularities of detecting user subjectivity through brain-computer interfaces, and present two emotion visualization techniques for touch and desktop interfaces. These visualizations offer real-time or post-task feedback about the users' affective states, both in single-user and collaborative settings, thus increasing the emotional self-awareness and the awareness of other users' emotions. For supporting collaborative interaction, a novel design for tabletop tangibles is described together with a set of specifically developed interactions for supporting tabletop collaboration. These ring-shaped tangibles minimize occlusion, support touch interaction, can act as interaction lenses, and describe logical operations through nesting operations. The visualization and the two UCCV techniques are each evaluated individually capturing a set of advantages and limitations of each approach. Additionally, the collaborative visualization supported by the two UCCV techniques is also collectively evaluated in three user studies that offer insight into the specifics of interpersonal interaction and task transition in collaborative visualization. The results show that the proposed collaboration support techniques do not only improve the efficiency of the visualization, but also help maintain the collaboration process and aid a balanced social interaction.

Keywords: collaborative information visualization, computer-supported cooperative work (CSCW), user-centered design, emotion visualization, affective user interface, tabletop, touch surfaces, nestable tangibles, subjective evaluation.

Acknowledgements

The completion of this thesis would not have been possible without the constant encouragement and support of my two supervisors, Prof. Dr. Achim Ebert and Prof. Dr. Andreas Kerren. Their guidance during my doctoral studies has been invaluable and their inspiration aided me in further defining my way in academia. I would also like to express my deep appreciation to Prof. Dr. Hans Hagen, through whose advice and comments I have obtained vital insights into my research as well as academia in general.

I would like to thank the colleagues and students I worked with at the University of Kaiserslautern and the Linnaeus University. Working with you kept me motivated through every step of the way, and I hope I managed to learn something from each one of you. I am especially thankful to Peter-Scott Olech, Sebastian Thelen, Christopher Weber and Sebastian Petsch for the many discussions and the fruitful collaboration on multiple research projects. Similarly, I would like to thank Ilir Jusufi, Björn Zimmer and Kostiantyn Kucher for their helpful comments throughout the years.

Special thanks to my close friends and colleagues Valentina Morar, Stefka Tyanova, Anuj Sehgal and Orsolya Eموke Sehgal. Your examples guided me through the rough waters that I have encountered during my doctoral studies. And when guidance alone would not suffice, you were there to push me onwards and get me moving again. Additionally, I would like to express my gratitude to Mady Gruys and Roger Daneker for their assistance in the official aspects of my studies as well as all their kind words.

This thesis would have been much more difficult to complete without all of the friends that stood by my side and offered their unconditional support. Most importantly, I would like to wholeheartedly thank my parents and my brother for their trust and endless support. None of this would have been possible without the strength that they have given me through their confidence and love.

Table of Contents

Abstract

Acknowledgments

Table of Contents

List of Figures

List of Publications

1	Introduction	1
1.1	Aims of this Thesis	6
1.2	Research Questions and Goals	8
1.3	Overview and Contribution	9
2	Related Research	13
2.1	Collaborative Visualization on Shared Multi-touch Displays	14
2.2	Emotions and Their Role in Collaboration	18
2.3	Visualizing the Web and User Online Browsing	31
3	Visualizing Parallel Web Browsing Behavior	35
3.1	Motivation	36
3.2	Requirement Analysis	38
3.3	WebComets: Tab-oriented Browser History Visualization	40
3.4	Summary	58
4	Emotions in Collaborative Settings: Detection and Awareness	61
4.1	Detecting Emotions with a Commercial EEG Headset	64
4.2	Detecting Insight through User Emotions	81

4.3	Emotion Visualization on Multi-touch Displays	96
4.4	Emotion Visualization on Desktop Interfaces	114
4.5	Summary	131
5	Collaborative Interaction with Parallel Browsing Histories	133
5.1	Nestable Tangibles for Collaborative Interaction on Tabletops	135
5.2	Supporting User-Centered Collaborative Visualization	151
5.3	Summary	168
6	Conclusions and Outlook	171
6.1	Discussion	171
6.2	Conclusions	177
6.3	Future Work	181
	Bibliography	183

List of Figures

1.1	Applegate’s place-time matrix of collaboration	3
1.2	Multiple users collaborating on a tabletop visualization supporting various input techniques	7
2.1	Russell’s circumplex model of affect	22
3.1	Multiple tabs open in the same web browser window	36
3.2	WebComets visualization of the parallel browsing histories of two users	41
3.3	Conceptual representation of a tab timeline	43
3.4	Conceptual representation of the tab hierarchy	43
3.5	Circular glyph representation of a visited web page	44
3.6	The list of supported categories together with the corresponding colors for the circular encodings	46
3.7	Example of a web page comet from the category “unknown”	47
3.8	Temporal zoom along the horizontal time axis	49
3.9	Connections between the selected web pages and other glyphs highlighted through continuous curved lines	49
3.10	Information box presenting details about the selected web pages	51
3.11	Collapsed view	52
3.12	The motif search window helps analysts construct, save and search for custom navigation patterns	53
3.13	Example of logical combinations of custom motifs and their corresponding search results in a parallel browsing log	54
3.14	WebComets visualization of multiple parallel browsing sessions	57
4.1	Structure and logical flow of the sections in Chapter 4	62
4.2	The Emotiv EPOC EEG headset (images courtesy of Emotiv)	65

4.3	The 10-20 electrode distribution system and the locations of the EPOC's electrodes	66
4.4	Percentage of correctly recognized facial expressions	70
4.5	Sample video frame from the user calmness measurements	71
4.6	Average difference between the EPOC device output and the questionnaire results for the emotional evaluation	73
4.7	Average difference between the EPOC device output and the questionnaire results for the two scenarios	75
4.8	Percentage of reduction for the false positives when considering groups of facial expressions	79
4.9	Representation of the eight-coin problem	87
4.10	Example of a matchstick arithmetic problem	87
4.11	Measured emotions with the EPOC headset in the presence of insight	88
4.12	Measured emotions with the EPOC headset in the absence of insight	89
4.13	ManyEyes map visualization employed in our experiments	90
4.14	Correlation between the number of insights and the instances where frustration, excitement and frustration-excitement pairs were detected	91
4.15	Correlation between the insights generated by the participants and the emotional responses detected by the EEG headset	93
4.16	Russell's circumplex model of affect extended by the visual metaphors employed by EmotionPrints	99
4.17	Computing the outline of the EmotionPrints halo	101
4.18	Example of EmotionPrints dissipating and being overwritten by newer touch instances	102
4.19	Histogram representation of touch events and their associated arousal-valence values for two users	104
4.20	EmotionPrints histogram can be displayed with constant time intervals or can be compressed to eliminate time intervals where the user did not execute any touch events	105
4.21	Screenshot of the eSoccer multi-touch game supporting up to four users	107
4.22	Two players interacting with the eSoccer game on a tabletop while their emotional cues are being interpreted via BCI and represented through EmotionPrints	108
4.23	User interacting with the browser history visualization on a tabletop .	109
4.24	User playing on a tabletop while wearing the BCI headset	110

4.25	Computational path for the EmotionPrints representations	111
4.26	Emotion data acquisition and representation loop	119
4.27	Different representations of EmotionScents	121
4.28	The emotion scent representation	121
4.29	EmotionScents representation applied for a combo box control	122
4.30	EmotionScents represented for slider widgets	123
4.31	EmotionScents encoding the user arousal during the interaction with the SwingSet demo interface	124
4.32	EmotionScents displayed on the interface of a simple Java IDE	126
4.33	EmotionScents-enhanced visualization for a dataset obtained from the ManyEyes website	127
5.1	TangibleRings in a map-based application, each ring controls different information layers	136
5.2	Prototyping the TangibleRings	140
5.3	The detection of the correct pairs of ring markers	141
5.4	Computation of ring orientation	143
5.5	Ring detection using custom markers for single and concentric rings	144
5.6	Interactions supported by the TangibleRings	146
5.7	Map-based application that supports TangibleRings and runs on a tabletop	147
5.8	Combining layers with the nested tangibles	148
5.9	An example of TangibleRings enabled interaction: rotation to control a view attribute, zoom mode, locked view	149
5.10	WebComets Touch and TangibleRings being employed on a tabletop display	152
5.11	Window presenting a user's browsing session	153
5.12	The constraints for a ring can be managed by accessing two property windows through the ring menu	154
5.13	Two nested TangibleRings in WebComets Touch visualization	156
5.14	Sum of multiple sine functions and the GAT representation for their mapping to a circle contour	159
5.15	Users working collaboratively on the WebComets tabletop application while wearing EEG headsets	160

5.16 Representation of the group affective tone cue during the interaction with the tabletop display 161

5.17 Representation of the valence fluctuation and the GAT for the three participants from one of the test groups 166

6.1 Types of interaction and awareness in user-centered collaborative visualization 174

List of Publications

This thesis includes ideas and materials from the following publications:

1. Daniel Cernea, Christopher Weber, Achim Ebert, and Andreas Kerren. Emotion-prints: Interaction-driven emotion visualization on multi-touch interfaces. In *Proceedings of the SPIE 2015 Conference on Visualization and Data Analysis (VDA '15)*, volume 9397, Burlingame, CA, USA, 2015. IS&T/SPIE (to appear). Materials appear in Chapter 4.
2. Daniel Cernea, Achim Ebert, and Andreas Kerren. Visualizing group affective tone in collaborative scenarios. In *Proceedings of the Eurographics Conference on Visualization (EuroVis '14)*, Poster Abstract, Swansea, Wales, UK, 2014. Materials appear in Chapter 5.
3. Daniel Cernea, Igor Truderung, Andreas Kerren, and Achim Ebert. An interactive visualization for tabbed browsing behavior analysis. *Computer Vision, Imaging and Computer Graphics – Theory and Applications, Communications in Computer and Information Science*, 458:1–16, 2014. Materials appear in Chapter 3.
4. Daniel Cernea, Igor Truderung, Andreas Kerren, and Achim Ebert. Web-Comets: A tab-oriented approach for browser history visualization. In *Proceedings of the International Conference on Information Visualization Theory and Applications (IVAPP '13)*, pages 439–450, Barcelona, Spain, 2013. SciTePress. Materials appear in Chapter 3.

5. Daniel Cernea, Christopher Weber, Achim Ebert, and Andreas Kerren. Emotion Scents – a method of representing user emotions on GUI widgets. In *Proceedings of the SPIE 2013 Conference on Visualization and Data Analysis (VDA '13)*, volume 8654, Burlingame, CA, USA, 2013. IS&T/SPIE. Materials appear in Chapter 4.
6. Achim Ebert, Christopher Weber, Daniel Cernea, and Sebastian Petsch. TangibleRings: Nestable circular tangibles. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI '13)*, pages 1617–1622, Paris, France, 2013. Materials appear in Chapter 5.
7. Daniel Cernea, Peter-Scott Olech, Achim Ebert, and Andreas Kerren. Measuring subjectivity – supporting evaluations with the Emotiv EPOC neuroheadset. *Künstliche Intelligenz – KI Journal, Special Issue on Human-Computer Interaction*, 26(2):177–182, 2012. Materials appear in Chapter 4.
8. Daniel Cernea, Achim Ebert, and Andreas Kerren. Detecting insight and emotion in visualization applications with a commercial EEG headset. In *Proceedings of the SIGRAD 2011 Conference on Evaluations of Graphics and Visualization – Efficiency, Usefulness, Accessibility, Usability*, pages 53–60, Stockholm, Sweden, 2011. Linköping University Electronic Press. Materials appear in Chapter 4.
9. Daniel Cernea, Peter-Scott Olech, Achim Ebert, and Andreas Kerren. EEG-based measurement of subjective parameters in evaluations. In *Proceedings of the 14th International Conference on Human-Computer Interaction (HCI '11)*, volume 174 of CCIS, pages 279–283, Orlando, Florida, USA, 2011. Springer. Materials appear in Chapter 4.

The following publications have not been included in this thesis:

10. Daniel Cernea, Christopher Weber, Andreas Kerren, and Achim Ebert. Group affective tone awareness and regulation through virtual agents. In *Proceedings of the Affective Agents Workshop at the 14th International Conference on Intelligent Virtual Agents (IVA '14)*, Boston, MA, USA, 2014. Springer.
11. Daniel Cernea, Achim Ebert, and Andreas Kerren. A study of emotion-triggered adaptation methods for interactive visualization. In *Proceedings of the 1st International Workshop on User-Adaptive Visualization (WUAV) at the 21st Conference on User Modeling, Adaptation and Personalization (UMAP '13)*, pages 9–16, Rome, Italy, 2013. CEUR-WS.
12. Daniel Cernea, Simone Mora, Alfredo Perez, Achim Ebert, Andreas Kerren, Monica Divitini, Didac Gil de La Iglesia, and Nuno Otero. Tangible and wearable user interfaces for supporting collaboration among emergency workers. In *Proceedings of the 18th CRIWG Conference on Collaboration and Technology (CRIWG '12)*, volume 7493 of LNCS, pages 192–199, Duisburg, Germany, 2012. Springer.
13. Daniel Cernea, Peter-Scott Olech, Achim Ebert, and Andreas Kerren. Controlling in-vehicle systems with a commercial EEG headset: Performance and cognitive load. In *Visualization of Large and Unstructured Data Sets – Applications in Geospatial Planning, Modeling and Engineering (VLUDS '12)*, pages 113–122, Schloss Dagstuhl, Leibniz-Zentrum für Informatik, 2012. OpenAccess Series in Informatics (OASICs).
14. Peter-Scott Olech, Daniel Cernea, Helge Meyer, Sebastian Schöffel, and Achim Ebert. Digital interactive public pinboards for disaster and crisis management – concept and prototype design. In *Proceedings of the 2012 International Conference on Information and Knowledge Engineering (IKE '12) at the 2012 World Congress in Computer Science, Computer Engineering, and Applied Computing (WorldComp '12), IKE2460*, Las Vegas, USA, 2012. CSREA Press.
15. Anuj Sehgal, Daniel Cernea, and Milena Makaveeva. Pose estimation and trajectory derivation from underwater imagery. In *Proceedings of the MTS/IEEE Conference Oceans 2012*, pages 1–5, Yeosu, South Korea, 2012. IEEE Computer Society Press.

16. Niklas Elmqvist, Andrew Vande Moere, Hans-Christian Jetter, Daniel Cernea, Harald Reiterer, and T. J. Jankun-Kelly. Fluid interaction for information visualization. *Information Visualization Journal (IVS), Special Issue on Information Visualization: State of the Field and New Research Directions*, 10(4):327–340, 2011.
17. Petra Isenberg, Niklas Elmqvist, Jean Scholtz, Daniel Cernea, Kwan-Liu Ma, and Hans Hagen. Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization Journal (IVS), Special Issue on Information Visualization: State of the Field and New Research Directions*, 10(4):310–326, 2011.
18. Daniel Cernea, Achim Ebert, Andreas Kerren, and Valentina Morar. R3 – un dispozitiv de intrare configurabil pentru interactiunea libera in spatiu (R3 – a configurable input device for free-space interaction). In *Proceedings of the 7th National Conference on Human-Computer Interaction (RoCHI '10)*, volume 3, pages 45–50, Bucharest, Romania, 2010. Matrix Rom.
19. Peter-Scott Olech, Daniel Cernea, Sebastian Thelen, Achim Ebert, Andreas Kerren, and Hans Hagen. V.I.P. – supporting digital earth ideas through visualization, interaction and presentation screens. In *Proceedings of the 7th Taipei International Digital Earth Symposium (TIDES '10)*, pages 36–49, Taipei, Taiwan, 2010.
20. Anuj Sehgal and Daniel Cernea. A multi-AUV missions simulation framework for the USARSim robotics simulator. In *Proceedings of the 18th IEEE Mediterranean Conference on Control and Automation (MED '10)*, pages 1188–1193, Marrakech, Morocco, 2010. IEEE Computer Society Press.
21. Anuj Sehgal, Daniel Cernea, and Andreas Birk. Modeling underwater acoustic communications for multi-robot missions in a robotics simulator. In *Proceedings of the IEEE Oceans 2010 Asia-Pacific*, pages 1–6, Sydney, Australia, 2010. IEEE Computer Society Press.
22. Anuj Sehgal, Daniel Cernea, and Andreas Birk. Simulating underwater acoustic communications in a high fidelity robotics simulator. In *Proceedings of the 7th IFAC Symposium on Intelligent Autonomous Vehicles (IAV '10)*, volume 7, pages 587–592, Lecce, Italy, 2010. Elsevier Science Ltd.

23. Anuj Sehgal, Daniel Cernea, and Milena Makaveeva. Real-time scale invariant 3D range point cloud registration. In *Proceedings of the International Conference on Image Analysis and Recognition (ICIAR '10)*, volume 6111 of LNCS, pages 220–229, Povoá de Varzim, Portugal, 2010. Springer.
24. Sebastian Thelen, Daniel Cernea, Peter-Scott Olech, Andreas Kerren, and Achim Ebert. D.I.P. – A digital interactive pinboard with support for smart device interaction. In *Proceedings of the IASTED International Conference on Portable Lifestyle Devices (PLD '10)*, Marina Del Rey, USA, 2010. ACTA Press.

Chapter 1

Introduction

The amount and complexity of digital data has been increasing exponentially for the last couple of decades. And while the storage technologies manage to retain this entire universe of ones and zeros, the same cannot be said about our ability to transform this data into knowledge and insight. Thus, data needs to be explored, analyzed and interpreted in order to transform it into *information*, a meta-view of the data that uses relationships to obtain higher level concepts. In other words, information is something we can easily understand and process, basically “differences that make a difference” [16]. At the same time, information can only capture various pieces of the puzzle, a puzzle that needs to be solved through patterns and comparisons in order to reach the final goal: *knowledge*.

While the complexity and size of the datasets is increasingly often addressed through automatic computational methods from fields like artificial intelligence, machine learning or computer vision, there are still domains and levels of complexity that result in some datasets requiring the experience and intuition of a human expert for extracting higher level knowledge. Humans have been involved in such analysis and knowledge extraction processes for centuries, however never on a scale similar to the one enabled today by the explosion of digital data. As a result, often large datasets are being explored and analyzed in powerful visualization systems that allow users to recognize patterns in large amounts of data and gather insights about the underlying information. This combination of novel technologies, intuitive representations and flexible interaction techniques allows users to actually see with their physical eyes what they previously could only perceive in their imagination, in best case [84].

A single person has a certain education, experience, intuition and cultural background, each of which can become both an asset and a limitation when exploring large amounts of data visually. To overcome this, “visualization must be a collaborative activity” [315]. As such, visualizations have been designed that explicitly and implicitly support the collaboration of multiple users to combine their experience and analytical power, as well as build upon each other’s knowledge and potentially reach deeper and more valuable insights. As further stated in [279], “the expertise to analyze and make informed decisions about these information-rich datasets is often best accomplished by a team”. Yet, visualizations that enable user collaboration do not simply allow multiple users to share a workspace, but also offer the means for supporting the dynamic process of social interaction and data manipulation.

Imagine a group of doctors preparing for a surgery by examining the patient’s medical record on a large multi-touch display. The visualization system being executed on the display does not only represent all the data and allow participants to interact with it. Ideally, a collaborative visualization enhances the way the doctors would communicate and interact, as well as support the exchange and manipulation of the patient’s medical information. Noticeably, this is a highly complex issue with many factors that need to be considered, but with the aim that “understanding these collaboration-centered socio-technical systems could accelerate their adoption and raise their benefits” [259].

Before diving into the various challenges that collaborative visualization faces and the subset of these challenges that we plan to address in this thesis, we need to take a closer look at the concept behind collaborative visualization. While the reasoning and advantages behind both collaboration and visualization have been briefly highlighted, we now have to address the concept of collaborative visualization specifically. Among the number of potential definitions discussed in [124], the one proposed by Isenberg et al. offers the most far-reaching and broad view on the field, by stating that:

Collaborative visualization is the shared use of computer-supported, (interactive,) visual representations of data by more than one person with the common goal of contribution to joint information processing activities.

When talking about computer-supported collaboration, we also need to highlight the definition of computer supported collaborative work (CSCW) as described in [15]:

Computer supported collaborative work should be conceived as an endeavor to understand the nature and characteristics of cooperative work with the objective of designing adequate computer-based technologies.

These definitions offer us a frame of reference, where the distinct elements of a successful collaborative visualization are formed by the users, the computer system (i.e., the domain data, the representation, and the interaction) and the goals. These elements need to be considered both individually and collectively when designing a collaborative visualization technique, as we will highlight later in this section.

	Co-located	Distributed	
	Single displays Large screens Tabletops	Virtual worlds Shared screens Online visualizations	Synchronous
	Single displays Large public displays Pinboards	Online visualizations Virtual pinboards Wiki systems	Asynchronous

Figure 1.1: Applegate’s place-time matrix of collaboration [7]. The four cells of the matrix highlight potential collaboration scenarios in the context of visualization. Our focus in this thesis goes towards the first cell, namely the co-located synchronous collaboration.

Moreover, collaborative visualization can be further described by inspecting the space-time matrix classification [7] that basically distributes collaboration based on the spatial location from where the users are interacting (distributed or co-located) and the moment in time when they are collaborating (synchronous or asynchronous).

Figure 1.1 presents Applegate’s space-time matrix of collaboration and specifies a couple of potential collaboration scenarios. Note that a system does not have to fall exclusively inside one of the four categories, e.g., online visualization websites where users can cooperate both in real-time and asynchronously.

Collaborative visualization is however a two edged sword—it promises an increased rate of knowledge gain as well as better solutions, but it is also challenging to design due to the inherent multifaceted complexity of the process [279]. Involving elements of computer graphics, perception, software development, interaction, cognitive and social psychology, etc., collaborative visualizations have to consider and sustain all these aspects in order to achieve the final goal.

Context is a keyword when talking about collaborative visualization. People are not mere machines that can focus on a task, analyze and extract information. The human condition involves elements of social interaction and close coupling with the physical environment. As such, people involved in collaborative scenarios are influenced by a set of contextual features, both in terms of external factors and internal experiences. Moreover, encompassing this topic under terms like situational or data awareness does not cover the entire space of the experience [82, 103, 260]. Certainly, awareness of other users’ actions and changes to the data does influence the entire analysis and decision making process. However, the contextual dimension should not be reduced to data-specific terms, as subjective user experiences and interpersonal interaction can also heavily influence the collaboration process. Users “need to be ‘aware’ of each other in such an environment in terms of their intentions, general feelings, influence on the shared workspace, etc.” [166].

Considering the multi-dimensional contextual information that our research is trying to address, we hereby propose the concept of *user-centered collaborative visualization*. The term is inspired by the concept of *user-centered design* (UCD), defined in [270] as:

An approach to user interface design and development that views the knowledge about intended users of a system as a central concern, including, for example, knowledge about user’s abilities and needs, their task(s), and the environment(s) within which they work. These users would also be actively involved in the design process.

This UCD definition, as well as corresponding definitions captured by ISO standards (ISO 13407 and ISO TR 18529), highlight the importance of fine-tuning the system or visualization for the specific abilities and needs of users, thus minimizing their efforts to adapt to the novel technology and maximizing their focus and productivity in the context of the tasks. *But, can we talk about a user-centered design in collaborative visualization?* We hereby propose the following definition for the concept of user-centered collaborative visualization:

User-centered collaborative visualization is the shared use of computer-supported, interactive, visual representations of data that considers knowledge about the abilities and needs of both the involved users and the group as such, their task(s), and the environment(s) within which they work, in order to capture vital contextual information and support the process of completing the user group's common goal of contribution to joint information processing activities.

In other words, a user-centered collaborative visualization (UCCV) is a user-centered design process in the context of collaborative visualization tasks, where the abilities, needs, tasks and environments of all the individuals and the resulting group are considered. While similar to the concept of human-centered visualization environments [140], UCCV incorporates additional considerations related to the social dimension of interpersonal interaction as well as the idea of group dynamics as a distinct entity. As such, the corresponding visualization techniques need to consider both the goals of individuals as well as the goal of the entire team. More importantly however, a UCCV system does not only support each user in his/her activity, but also the entire group as an independent entity with requirements that allow it to function efficiently and abilities that exceed the ones of all the individuals in the group (i.e., the whole is greater than the sum of its parts). Similarly to UCD, UCCV requires an active involvement of both users and the group, in order to obtain a balance between the support offered to the individuals and to the group as a whole entity.

1.1 Aims of this Thesis

As the exploration space for collaborative visualization is clearly vast, in this thesis we aim at supporting collaborative visualization in co-located synchronous tasks around tabletop displays. In this context, we consider the following eight design guidelines proposed in [251] which are specifically aimed at supporting effective tabletop collaboration:

1. support interpersonal interaction,
2. support fluid transitions between activities,
3. support transitions between personal and group work,
4. support transitions between tabletop collaboration and external work,
5. support the use of physical objects,
6. provide shared access to physical and digital objects,
7. consider the appropriate arrangements of users, and
8. support simultaneous user actions.

Based on these guidelines, the research presented in this thesis highlights collaborative visualization topics related to social interaction (i.e., how can we support the inter-user collaboration) and system interaction (i.e., how can we support the man-machine interaction), thus addressing the first three guidelines from the previous list. For the social aspect, we considered the interpersonal interaction by supporting and exploring the impact of user emotional awareness in tabletop collaboration. While nowadays cognitive psychologists agree that cognitive and affective features are heavily interconnected [257], there are still a limited number of systems, including in the realm of CSCW, that take both mental processes and affect into consideration.

Still, emotional states are not only omnipresent in humans, but also an integral part of reasoning and communication [64,182,238]. User emotional states have been shown to have an effect on creativity [89], motivation [60] and problem solving [64]. More importantly for the context of collaborative visualization, emotions have been connected to performance in visual tasks [102, 165], learning [152], and decision making [166].

As such, researchers now start to recognize that the major challenges of collaboration might be social and not technical [59, 109], and that these challenges—while globally more difficult—might improve the performance and nature of collaboration in significant ways.



Figure 1.2: Multiple users collaborating on a tabletop visualization [46]. Their collaboration and interaction with the system can be supported through touch events, gestures, tangibles, and other interaction metaphors.

Secondly, for the system interaction, we considered potential techniques for actively aiding flexible and user-centered collaboration in the context of fluid transition between various activities. As highlighted in Figure 1.2, tabletop interaction breaks the standardized approach of the desktop computer and enables a flexible interaction based on various techniques for detecting touch events, gestures and even a wide range of objects. Whatever the concept behind the interaction, one of the main focus points when developing a multi-touch interface is to seamlessly integrate the various interaction concepts and metaphors [175]. In this context, tangibles—physical objects whose presence and manipulation can be detected by the tabletop system—present a high potential for reusing user mental models.

More specifically, well-conceived tangibles have a design that allows users to quickly deduce their visual metaphor (miniature houses or cars that can be positioned on the tabletop surface [136, 177, 286, 289]) and/or their function (a knob-like object, a checker piece [110, 209, 225, 308]), by drawing on comparable features to objects or concepts from the real world. Given the wide range of metaphors that can be employed for interactive scenarios, our aim was to leverage the experience and dexterity of users with manipulating physical objects and building on these common real world interaction models.

1.2 Research Questions and Goals

The topics that are being inspected in this thesis focus on supporting the collaboration visualization process through a user-centered approach. More precisely, one of the major problems in collaborative visualization is constituted by the fragility and complexity of the interpersonal interaction and cooperation. Often, collaborative visual solutions focus merely on supporting the multi-user functionality of an application, while at the same time neglecting to develop or extend interaction techniques to satisfy the requirements of a truly collaborative system. Moreover, the social dynamics of a collaborative session are often neglected or ignored. Yet, collaboration tends to have a highly complex interplay of social and psychological aspects for each participant, aspects that are rarely considered and which have the potential to *make or break* the entire collaboration.

The main **goal** of this research is to explore and develop novel techniques for supporting collaborative visualization through a user-centered approach that considers both conscious and subconscious features and their relevance to collaborative scenarios. Specifically, the criteria for fulfilling our goal are defined as:

- A. Develop a novel visualization in order to test and evaluate our proposed UCCV techniques in a realistic (collaborative) visualization scenario.
- B. Create awareness of user subjective aspects in collaborative sessions in order to improve group cooperation, communication and overall social dynamics.
- C. Design custom interaction techniques that “naturally” support collaborative visualization and address issues like privacy, private and shared space, etc.

- D. Evaluate the proposed interaction and visualization techniques both through conventional and subjective measures in order to capture the quantitative and qualitative aspects of the proposed solutions.

Throughout this thesis, these criteria are mainly tackled in the context of co-located synchronous collaborative information visualization, with individual solutions that can be extended to the other quadrants of Applegate’s place-time matrix.

1.3 Overview and Contribution

This thesis covers a wide range of research topics driven by the goal criteria highlighted in Section 1.2. Each of these criteria will be addressed individually, highlighting aspects like motivation, implementation, functionality, evaluation, use cases and more. A final chapter will connect all these modules in a single collaborative visualization system that will be further analyzed from the perspective of our initial goal.

In Chapter 2, we give an introduction to collaborative visualization and collaborative visualization on tabletop displays, specifically. We further present relevant research in the areas closely connected to concepts and techniques introduced in this dissertation: emotion visualization and awareness, subjectivity measurements, touch and tangible interaction, and user parallel web browsing.

In Chapter 3, we introduce WebComets, an interactive visualization technique for analyzing and comparing user parallel (i.e., tab and window-based) web browsing behavior in multi-user multi-session online histories. Besides the representation and its interactive features, the visualization also supports advanced filtering based on both content and context information, allowing users to detect and compare highly complex patterns in the parallel browsing logs. The development of the concept and visualization are preceded by a requirement analysis focusing on the optimal approach towards a scalable visual representation for long-term parallel browsing logs. An evaluation, a discussion and a use case further support the abilities of our visualization tool to aid the analysis of parallel online browser histories. The main contributions of this chapter include:

1. A novel interactive visualization for representing and analyzing parallel web browsing behavior, one of the first visualization tools aimed specifically at analyzing and finding patterns in tabbed browsing data.
2. A motif-enabled search system aimed specifically at detecting and comparing temporal and semantic patterns—based on content and context information—in the online browsing graph.

In Chapter 4, we explore the potential of lightweight EEG devices to correctly recognize and classify user emotions in order to use this information for enhancing affect awareness through emotion visualization and evaluating visualization through subjectivity measurements. For this purpose, we compare the emotion classification obtained by using an EEG headset and its software framework to video logs and user self-reports. The results of the validation session are highlighted, and a discussion captures the advantages and limitations of our approach. Moreover, in a new set of experiments we identify a strong correlation between specific user emotions—e.g., frustration and excitement—and moments of insight, further reinforcing the importance of emotion recognition technology in the context of evaluating visualization techniques through user subjectivity.

Next, we describe the concept behind an emotion visualization framework for representing user emotional states in multi-touch (EmotionPrints) and desktop (EmotionScents) environments. A set of design guidelines is highlighted for real-time, in-place visualization of user emotions, i.e., the visualization of the users’ affective state at the moment and the location where an action is executed in the interface. Further, a set of examples capture the potential of our framework in the context of emotional self-awareness, the awareness of team member emotions in collaborative settings, and interface evaluation. Subsequent evaluations and discussions address the framework’s ability to convey emotional states consistently and accurately. The scientific contributions of this chapter include:

1. A study capturing the potential and accuracy of using lightweight EEG technology to capture user affective states, and employ this information in the context of subjectivity-based evaluation.
2. A study highlighting the correlation between specific emotional states (e.g., frustration, excitement) and moments of insight, as well as the potential of using lightweight EEG devices for detecting moments of insight in visualization.

3. A visualization framework for representing user emotional states on both multi-touch and desktop (i.e., widget-based) interfaces, irrespective of the technology used for emotion acquisition.

In Chapter 5, we introduce a novel ring-shaped tangible for supporting intuitive and collaborative interaction on tabletop displays. We describe the concept and design of the TangibleRings as well as highlight the functionality in terms of individual and shared workspace use. In focus is the ability to combine or nest the TangibleRings, allowing users to generate complex filters, views and queries in the form of a physical lens. A study and a set of application examples further capture the value of the TangibleRings metaphor.

Next, we present a modified version of the WebComets visualization described in Chapter 3 that is specifically designed for tabletop collaboration and is augmented in terms of interaction by the TangibleRings. We evaluate our tabletop visualization, the corresponding interaction techniques and the collaboration support through both qualitative and quantitative measures, by again employing EEG technology as described in Section 4.1. We conclude by discussing our findings. The main contributions of this chapter include:

1. A novel concept for nestable passive tangibles employed on tabletop displays in co-located collaborative settings.
2. A collaborative, interactive visualization for representing and analyzing parallel web browsing behavior on tabletop displays.
3. A study on the value and potential of an evaluation based on user subjectivity in a collaborative visualization system.

Finally, in Chapter 6 we present the conclusions of our research and address potential future extensions for each of the major contributions highlighted above.

Chapter 2

Related Research

Collaborative visualization is still a relatively new area of research, defined by the intersection of computer-supported cooperative work (CSCW) and visualization. In the initial years, collaborative visualization was revolving mainly around large scientific datasets from fields like biology [25, 26, 87], medicine [98, 173, 269], chemistry [43, 176], geography and cartography [31, 169], etc. And while the beginnings of collaborative visualization are in the scientific community [33], nowadays many such collaborative techniques and representations have been specifically designed for “abstract, nonspatial data” [283] in the context of information visualization [124, 140].

As mentioned in the introductory chapter, collaborative visualization can be divided based on the location (co-located and distributed) and time (synchronous or asynchronous) of the collaboration [7]. Furthermore, in terms of the interaction that can take place in collaborative scenarios, visualizations support multiple interaction categories, including viewing the data, interacting and exploring the dataset, and creating new representations or sharing them with others [198].

Considering the large number of research publications on the topic of collaborative visualization, the current chapter focuses on highlighting only research that is more closely related to the topic of this thesis. As such, our attention goes towards user-centered collaborative visualization with an emphasis on issues related to collaborative information visualization in co-located synchronous interaction, involving viewing, exploration, creation and sharing of data. More specifically, the sections in this chapter highlight the state-of-the-art for: tangible objects for supporting collaborative tabletop interaction (Section 2.1); brain-computer interfaces (BCIs), affective computing

and emotion visualization with an accent on collaborative scenarios (Section 2.2); and the visualization of user online browsing behavior and patterns (Section 2.3). Note that an overview of the research in the context of the various cells in Applegate’s matrix (see Figure 1.1) can be obtained in [33, 100, 104, 124, 315].

While clearly addressing different areas of research, these topics are all connected by the topic of this thesis, namely user-centered collaborative visualization. On one hand, the user of novel interaction metaphors around the tabletop display and the representation of user emotions is meant to explore the their potential in supporting collaborative visualization sessions. On the other hand, the visualization of users’ online navigation represents the ideal candidate for an collaborative and interactive application, that not only an unsolved problem, but also offers us a backbone visualization for testing and evaluating the proposed UCCV techniques.

Note that the order of the sections in this chapter, sections that highlight the related research for the previously mentioned topics, is inverse to the order in which these subjects have been addressed in Chapters 3-5. This is motivated by the desire to introduce these concepts in a top-down fashion, starting with collaborative visualization, moving towards the specifics of collaboration and interaction on tabletop displays, and continuing with the importance of social interactions, specifically emotions, in these interpersonal group settings. Finally, as an application scenario for (collaborative) visualization, the state-of-the-art in online browsing and website navigation visualization is presented.

2.1 Collaborative Visualization on Shared Multi-touch Displays

Large displays are one of the main technologies that support collaborative visualization today. Representing mostly the segment of co-located synchronous collaboration, large screens allow multiple users to view and manipulate visualizations while sharing the same physical space. This is usually achieved through technologies like large high-resolution displays [144, 226, 228, 278], tabletop displays [126, 131, 228, 248, 249, 309] or multi-display environments [86, 87, 221, 299].

In this context, research has focused on a wide range of topics related to collaborative visualization and the corresponding user interaction: coordination and synchronization of interaction on a large multi-touch display [88, 125, 193, 195, 250, 281], orientation and user perception [22, 42, 154, 310], and awareness of changes in the data or representation [125, 141, 172, 276, 281].

Besides touch and gesture-based manipulation, another input approach in the context of these large displays is represented by the various interaction metaphors. One of the most famous ones in this sense is represented by the magic lens metaphor, which “filters are new a user interface tool that combine an arbitrarily-shaped region with an operator that changes the view of objects viewed through that region” [271]. Today, magic lenses are widely used in various visualizations systems that rely on tabletop displays. Their functionality allows multiple users to explore the same representation inside a local view that does not affect the tasks of the other users. Examples in this sense can be found both in the context of software lenses [21, 88, 137, 215] and hardware lenses [147, 263, 264]. While quite similar, the software solution usually offers more flexibility in terms of the interaction and a better integration in the visualization. On the other hand, lenses that are based on tangible objects offer the user both tactile feedback, precision and even advantages in terms of privacy.

2.1.1 Tangible Objects and Multi-touch Interaction

One way of enhancing collaborative interaction on tabletops is with the help of tangible objects. Tangibles are physical items that, when positioned on or around a tabletop, can extend the interaction capabilities of the users by augmenting the touch interaction, adding a third dimension to the display and/or supporting the collaboration process. This often results in a natural integration of virtual and generic interface elements, integration that is based on the extended experience of most users with interacting with objects on horizontal surfaces.

In terms of their functionality, tangibles can either be represented by every day objects (e.g., cups, notebooks) that are given additional meaning and capabilities with combined with the interaction capabilities of the tabletop, or custom artifacts specially designed for complementing touch-based interaction. At the same time, tangibles should be designed in such a manner that allows users to quickly distinguish the interactions

they support. Tangibles that employ a visual metaphor rely on their representation to convey the user what manipulation features they support (e.g., miniature cars, buildings, trees that can be positioned on the tabletop [136,175,286,289]), while those that employ interaction metaphors incorporate visual cues that communicate their functions to the users (e.g., knobs for turning or arrows for pointing [110,130,209,225,308]). Note that tangible user interfaces (TUIs) not related to tabletop displays and tangible objects are not considered in this section.

In recent years, there have been many approaches involving tangibles in different application areas [253], such as information visualization [61, 113], urban planning [118,289], emergency management [46], and entertainment [209,245]. A more technical approach is that of Bricks et al. [85], in which they introduce the concept of graspable user interfaces through the use of physical handles as controls. They used these handles to operate on top of a large horizontal display surface called “ActiveDesk”. Other activities and overviews on tangible user interfaces can be found in the works of Ishii [127, 128] or Underkoffler [288]. A more recent example can be found in [174, 175], where ColorTable is presented: a system used for urban planning that employs tangibles of different shapes and colors to represent a wide range of real-life objects. While most of these approaches are based on tangibles implemented through solid, opaque objects [175,264,289], there are a couple of systems that allow users to execute touch operations on the tangibles or reuse their area to highlight information [225,246,298].

With the Urp workbench [289], a system targeting urban planning and architectural purposes, Underkoffler et al. tried to minimize occlusion. Skeleton-like models of buildings are used to simulate lightning, wind, and shadow conditions during the design phase. In the work of Patten [208], Sensetables are introduced as physical objects that can be tracked electromagnetically on the surface of the tabletop. This approach allows data to be represented also on the tangibles, as a projector is positioned above the table in order to use both surfaces—tabletop and tangibles—as screen area. However, in such a projection-based approach, the users can easily add occlusion by moving their arms between the surface and the projector. At the same time, metaDESK [286] introduces tangible lenses that augment the displayed information and minimize occlusion, but without supporting touch interaction. Similarly, Völker et al. [295] describe the design of their transparent tangibles that are detected by capacitive touch displays in the same way as fingers. Note that capacitive sensors

are a technology that support the detection of conductive objects (e.g., user fingers) that are in the proximity or in contact with the display. While widely used in multi-touch displays, there are few tabletop displays that support capacitive technology due to the size limitations and costs of this approach. However, in the following we will also highlight a set of tangible objects designed for capacitive displays, but only if these tangibles are designed to rest on the surface of the touch screen, thus making their functionality similar to the one of tangibles in the context of tabletops.

The particular case of transparent and translucent tangibles is systematically explored in [36]. Besides a review of existing research in the area of tabletop tangibles, Büschel et al. offer a thorough view of the design space of transparent tangibles, their advantages and current limitations. Closely related, in [308] transparent tangibles of various shapes are employed as common input devices (e.g., knobs, keyboards, sliders). Jetter et al. [130] introduce another transparent tangible solution for supporting collaborative product search on tabletop displays. Besides some of the widely accepted manipulation methods, these facet tokens can also be connected to chains of logical criterions, thus formulating complex queries.

Baudisch et al. addressed a similar set of problems through their approach called Lumino [17]. Lumino employs cube-shaped tangibles filled with glass fiber bundles that traverse the objects from top to bottom. This allows users to stack multiple tangibles to build tridimensional structures, execute basic touch operations as well as perceive a low resolution image of the underlying screen area. Still, interaction possibilities are restricted to a simple touch/no-touch detection on the blocks. CapStones and ZebraWidgets [54] further extends Lumino's concept of structured transparency to capacitive tangibles consisting of multiple parts. This approach is only targeting devices with capacitive touch screens and cannot be used on tabletop computers that are based on diffuse illumination.

In terms of combining tangibles, Stackables [145] introduced a design for physical widgets that contain a display and two control wheels. With the help of the controls, users could adjust and select the facet values, which would also be visible on the display of the widget. At the same time, these facet values could be modified, stored, transferred and shared. Stackables could also be combined into vertical stacks, effectively forming more complex queries that would represent logical combinations of the

individual selections. However, the Stackables are designed as a self-sufficient tangible user interface and not as physical objects for augmenting tabletop interaction. Moreover, they only support a limited number of logical connectors, which limits the complexity of the queries that can be formed.

While many of the highlighted solutions address issues related to collaboration, occlusion, combined tangibles, interior touch detection, and cost, none of them manages to address all these tangible-related concerns at the same time. Our tangible solution, described in Section 5.1, offers an alternative to the presented approaches by addressing all of these issues and highlighting the advantages in the context of collaborative information visualization.

2.2 Emotions and Their Role in Collaboration

All of us experience emotions throughout our entire lives and in the most varied circumstances. Emotions influence our social interactions and interpersonal relations [59, 109] as well as everyday human processes like creativity [89], problem solving [64], motivation [32, 60], etc. And while emotions manage to spill into almost every aspect of our days, there are still some domains that are often ignoring human affective states and their effects.

One such context is given by our interactions with technology. Nowadays, we may work with specialized tools, communicate online, organize our lives with apps, have fun with games and experience an entire set of emotions while doing so, anything from joy to excitement and to frustration. Yet, while we do bring our emotional selves into the world of computing, this field has still very little in the way of supporting our emotion-escorted interaction. We communicate our subjective states online through emoticons [300], emojis [237] and “like” buttons, we smile in the face of an aesthetically pleasing interface but without any reaction from the system, we work together with our colleagues without knowing how they feel about our task or progress. As stated by Picard [214] in her work in the field of Affective Computing: “Basic affect recognition and expression are expected by humans in communication. However, computers today cannot even tell if you are pleased or displeased. They will scroll

screenfuls of information past you regardless of whether you are sitting forward eagerly in your seat, or have begun to emit loud snoring sounds”. All this makes our current technology seem dry, without the ability to capture the entire human experience and interaction, and center on the most important element of the equation: the *user*.

Focusing on our research, emotional states have also been linked to collaboration performance [166] as well as empathy [115,143]—a driving force for collaboration. At the other end of the spectrum, affective states have been associated with cognitive performance and visual judgement [102,165], two relevant attributes in working with visualization systems. As such, the measurement and integration of user emotional states into (collaborative) visualization solutions seems to represent the next step in terms of aiding the process of insight gathering and supporting collaborative analysis.

To consider this affective side of interpersonal interaction in technological settings, we need to be able to recognize or estimate user emotional states and employ these findings in the design, functionality and interaction-support of our digital systems. But before we focus on how we can use technology to detect emotions and how to convey them in single-user or collaborative settings, we need to clarify what emotions are and how they are classified. Certainly, subjective human experiences are inherently subtle and hard to define, resulting in publications that freely mix concepts like affect, emotion, mood, and sentiment. Furthermore, cognitive psychology offers a relatively wide array of definitions for all these concepts, leaving many degrees of freedom in selecting a view on the topic.

In the following, we try to offer a brief overview of the most widely accepted definitions and classifications of affective states. The notion of affective state covers a set of concepts, including *core affect*, *emotions*, *moods*, and *personality*. These concepts differ in multiple ways, like their triggers and their temporal persistency:

Core affects are defined as “a simple primitive non-reflective feeling most evident in mood and emotion but always available to consciousness” [232]. Core affect is thus constantly present in an individual, and it can be experienced either as part of an emotion or mood, or completely independently [231]. Furthermore, core affect can be linked to Russell’s circumplex model of affect, detailed later in this section. Some elementary building blocks of emotion that are included in the core affect category include concepts like pleasure-displeasure and energetic-tired.

Emotions are described as a medium term affective state that is characterized by their appearance as a response to an external or internal stimulus [241, 242], represented by a person, an event, a memory, an image, a scent, etc. Thus, emotions are not bounded by time or reality, as emotions can be generated also by imaginary experiences or by events from the past. Examples of emotional states include love, fear, anger, sadness, etc.

Moods are generally not elicited by a concrete stimulus, thus having a more diffuse nature. While the origin of a mood is rarely known, it is also an affective state that remains active for longer periods of time than emotions [90].

Personality is a more persistent subjective aspect encoding attitudes towards a concept or object [8, 192]. It can be defined in terms of the “Big Five” personality factors [97]: extraversion, agreeableness, conscientiousness, neuroticism, and openness to experience. More importantly, “two of the Big Five core personality factors can almost be defined in terms of emotional responding” [257], further supporting the close connection between personality and emotional experiences.

All these subjective states are of particular importance in collaboration, visualization and interaction. Especially emotions have a high relevance for the topic of this thesis, as these can be linked to events or objects, thus offering feedback about the interaction cycle of a user with a system or other users.

In terms of classification, emotion theory has centered around two ways of grouping affective experiences. On the one hand, some theories focus on defining and analyzing emotions as a set of distinct states. Some of the more widely used theories include: Ekman’s theory of six basic emotions [79, 206] (disgust, anger, fear, joy, sadness, and surprise; also called the “big six”) and Plutchick’s theory of eight basic emotions [217] (disgust, anger, fear, sorrow, joy, acceptance, anticipation, and surprise). Emotions that do not fall into the category of basic ones are usually defined as combinations or variations of basic emotions [29].

One characteristic of the basic emotions is that they are, by their nature, easier to recognize with various technologies, as they would generate distinct functional patterns in the human body (e.g., brain activation patterns or physiological patterns). However, studies have shown that non-basic emotions like frustration, boredom and confusion can be more frequent and thus more useful in human-computer interaction scenarios [70].

On the other hand, there are theories of affect that focus on distributing and describing all emotional states through a set of dimensions [96]. While the number of possible considered dimension is variable, most widely accepted approaches focus on a 2D or 3D model: Russell’s circumplex model of affect (see Figure 2.1) encodes emotional states in a two-dimensional space defined by valence (positive-negative) and arousal (excited-calm) [230], while the three-dimensional model of Mehrabian incorporating the three axes of pleasure (valence), arousal and dominance (abbreviated as PAD) [185]. Note that all these affective theories are mainly considered in the context of universal emotions, where emotional states are independent of cultural aspects [78, 186, 217].

Throughout this thesis, Russell’s circumplex model has been selected as the emotional reference system. Its two-dimensional space positions a set of distinct affective states—28 affect items in [230], 191 affect items in [233]—in terms of valence-arousal dyads. The model follows the dimensional theory of emotions that states that different emotions can be considered as dissimilar only in terms of one or more distinct dimensions [156].

While other models are available [72,217], Russell’s model is used also in other affective systems in visualization [267]. Additionally, arousal has been found to be related to valence, as “arousal increases when valence extends towards extremely positive of negative values. But for low values of valence, regardless of the polarity, arousal is almost always low.” [96] This interconnection has the potential to further reinforce the correct interpretations of the detected emotions. Additionally, it is noteworthy that Ekman’s six basic emotions have a correspondence in Russell’s two-dimensional model, basically offering a correspondence between the two categories of models.

The following sections give an overview of the technologies and approaches that are being employed for detecting and visualizing user affective states. In Section 2.2.1, our focus will go towards brain-computer interfaces (BCI) and, more specifically, electroencephalograph (EEG) measurements. The state-of-the-art in lightweight, mobile EEG headsets is highlighted, as these devices are being used increasingly often for interpreting user affective states, while at the same time offering a non-invasive solution to subjectivity measurements. Furthermore, Section 2.2.2 gives an overview of the emotion visualization techniques with a potential of increasing emotional awareness and positively impact user experience. These can, in turn, affect user performance both in single-user and collaborative visualization scenarios, as discussed in Chapter 4.




Figure 2.1: Russell’s circumplex model of affect [230] encoding user emotions in terms of two dimensions: valence (positive and negative emotions) and arousal (high and low excitement).

2.2.1 Detecting User Emotions

While still rare in the computational community, emotion measurement technologies are being used increasingly often in a variety of commercial and research projects, focusing mostly on the context of affective interaction and evaluation. However, these technologies differ a lot in terms of the modality that they inspect and the emotional variables they can detect [69,243]. This is mainly due to the fact that these devices do not detect emotions directly, but instead, they estimate them based on a set of human responses that are associated to affective states. More concretely, technologies that estimate user emotions can be grouped into the following three categories:

- *Perception-based estimation*: includes all elements of human emotion expression and behavior, including facial expressions, voice tone and modulation, hand and body movements, etc.

- *Physiological estimation*: focuses around the subconscious responses of the human body (e.g., heart beat, blood pressure, brain activity, sweating, etc.) to certain emotions. By means of specific technologies—e.g., Electroencephalogram (EEG), Electrocardiogram (ECG), Electromyogram (EMG), Electrodermal Activity (EDA), etc.—these subconscious responses can be monitored and employed as indications of the presence of certain emotional state. These responses may involve the central nervous system (CNS), the neuroendocrine system (NES), and the autonomic nervous system (ANS).
- *Subjective feelings*: are self-reports of individuals about their perception of their own emotional states. This category of techniques is widely employed in user studies from psychology and human-computer interaction, and it is less technology-dependent than the previous two.

When comparing the self-reports to the device-based emotional readings, a set of advantages and limitations can be detected. The technology-assisted approaches do not rely on the user’s self-reports and are therefore less affected by subjective perception. At the same time, employing perception or physiological estimation of affective states means that elements like culture or spoken language should not influence the results, contrary to the self-reports. On the other side, device-enabled emotion measurements need to be trained or calibrated for the individual. Furthermore, the acquired signals can be prone to noise as well as affected by external events, adding a new level of uncertainty to the affective data set (e.g., measuring users’ pupil size under changing light conditions, measuring users’ blood pressure levels when they are angry, etc.).

One particularly relevant aspect for emotion detection is that each technique inspects a specific set of affective indicators, and thus is better suited to offer information about certain emotional states or dimensions. For example, discrete emotions are usually well perceived through facial expressions or self-reporting, arousal through blood volume pulse measurements or heart rate, and negative valence through facial electromyography (EMG) and cortisol measurements.

In terms of emotion detection, the focus of this thesis is on detecting user affective states through readings obtained by non-intrusive, mobile brain-computer interface (BCI) devices. A BCI device can be defined as “a system to provide computer applications with access to real-time information about cognitive state, on the basis of measured brain activity” [317]. In other words, BCI is an umbrella term for a cate-

gory of devices that can detect and record the brain activity of a user, independently of the physical or chemical properties that are being explored. As such, BCIs can be subdivided into a number of technologies, each of which offering a different set of advantages and disadvantages in terms of resolution, portability, side effects, and invasiveness. A couple of well-known technologies include functional magnetic resonance imaging (fMRI), functional near-infrared brain monitoring (fNIRS), positron emission tomography (PET), and electroencephalogram (EEG).

From all these, besides their application in the field of medicine, EEG devices are most often employed in interaction, evaluation and subjectivity measurements, mostly due to its low cost, portability and non-intrusiveness. For this same reasons, this thesis is focused on estimating user affective states based mainly on EEG readings. The idea behind electroencephalographic measurements was first introduced in 1912 by Vladimir Pravdich-Neminsky who measured the electrical signals produced by the activation of firing neurons in the cerebral cortex of dogs [220,275]. With a principle that barely changed in over a hundred years, EEG measurements rely on a number of electrodes positioned at key locations on the scalp, and which can capture and amplify the activity of the neurons situated in their proximity. This means that modern EEG devices are still composed of these three basic parts: sensors that are called electrodes and are placed on the scalp of the user at key locations; an electronic amplifier that multiplies the signal obtained from the various electrodes; and a processing unit that either directly extract signal features or forwards the data to another system.

Focusing on the electrodes, these are usually placed following the international 10-20 system that encodes key locations on the human scalp, highlighting at the same time the corresponding lobes for each sensor location (see Figure 4.3). One should note that the original or the extended 10-20 system (19 and 70 electrodes, respectively) do not require an exhaustive use of the standard locations, as the guideline can be employed with a smaller number of the electrodes positioned at a subset of positions. Further, an EEG usually requires wet sensors to improve the quality of the electrical signal from the scalp, but dry sensor solutions are being improved and used increasingly often in commercial EEG devices.

In terms of processing, the electrical signal obtained by EEG measurements is in the area of 5 to 100 μV , requiring a high quality amplifier to avoid introducing additional noise into the signal. At the same time, the actual information carried by the electrical signals is in their frequencies, which are usually categorized in five frequency bands: delta (0Hz to 4Hz), theta (4-8Hz), alpha (8-13Hz), beta (13-30Hz), and gamma (>30Hz). The alpha and beta bands are most often inspected in the context of interaction and subjectivity, as alpha waves encode a relaxed, conscious state, and beta waves an active, engaged one. Moreover, both alpha and beta waves have been linked to changes in arousal and valence levels. More precisely, various patterns in alpha and beta waves measured around the frontal lobe have been linked to arousal and attention levels [3, 66], while asymmetric hemispheric readings of the alpha levels in the frontal lobe have been shown to encode valence [67]. Note that an overview of current-day commercial EEG devices can be found in [146, 318].

BCI systems can be employed for three main types of tasks: control and interaction, evaluation and subjective measurements [200]. Recent research in Human-Computer Interaction has focused mostly on the first category of applications. Solutions have been proposed that provide people that have motoric disorders with interaction possibilities [37, 157, 244]. Ranky et al. [223] introduced the use of the EPOC headset as a controlling device for a robotic arm, while in [2] an EEG device is used to control a hand in virtual reality. In terms of specific interfaces, Campbell et al. [38] utilize the EPOC headset to interact with mobile phones (e.g., dialing a number via EEG). In [48] vehicle functions (e.g., radio volume, side windows) are controlled via EEG measurements.

As the third category focuses mostly on user cognitive and affective states, the concept of affective brain-computer interfaces or aBCI [199] was introduced in order to group the approaches capable to detect indicatives of emotional states. These indicatives can be interpreted both as discrete emotional states or as dimensional features of affect.

An emotion classification model that is valid for most aBCI devices is described in [129]. Once an EEG device is mounted on the user's head and the electrodes are receiving a sufficient signal, the acquisition process can start. As the recorded signals are also affected by small electrical currents from muscle and eye movements, in the next step, a filter is applied that removes the electrooculogram (EOG) and elec-

tromyogram (EMG) signals, while at the same time amplifying the EEG one. Note that some EEG systems do not filter out the EMG signals, as these can encode additional affective information through the activation units of the facial muscles. Next, relevant features are extracted from the electrical signal, an operation that usually focuses on the various wave bands. Similarly to the classification of facial expressions, the following stage applies a set of machine learning algorithms to classify the information into emotional states of the user. A more detailed view of the classification of user emotional states based on different metrics—i.e., event related desynchronization/resynchronization (ERD/ERS), P300 response, and Steady State Visual Evoked Potentials (SSVEP)—is given in [191].

Additionally, multiple research publications have focused on EEG signal classification algorithms for emotion recognition [116, 149, 188, 196, 213]. Closely connected to affective states, cognitive metrics like mental workload have also been measured through custom algorithms [6, 99]. The classification accuracy of affective states based on EEG readings varies depending on the number of emotions that are being detected and the classification algorithm. For a simple bipolar valence measurement, an accuracy of 71% to 85% has been reported [129, 149]. Moreover, for 2 to 8 emotional states, recognition rates between 50% and 90% have been shown [262]. However, compared to the performance of facial expression-based emotion estimation, reliable EEG classification is still relatively difficult to achieve [199]. One reason for this can be the incorrect use of the EEG equipment.

While there are clearly different approaches to interpreting user emotions, in this research we focused on employing an approach that is portable and non-intrusive, in order to ensure its applicability in the context of tabletop collaboration. Our EEG-based detection, that is described and validated in Section 4.1, allows users to freely move and interact around large displays, while their brain signals are transmitted and interpreted as affective states in real-time.

2.2.2 Emotion Visualization

While still a novel area of research, the visualization of emotions has been addressed in multiple contexts over the last decades. Emotional states have been extracted from a variety of sources for the purpose of awareness through visualization, sources that include language/text [152, 163, 181, 203, 303], audio and video recordings [142, 159, 202], etc. However, our focus in this thesis goes towards visualizing affective states interpreted directly from users, especially for the purpose of emotional awareness and reflexion. Similarly to the awareness of interaction, situation and data [23, 82, 103, 260, 291], emotional awareness can affect the quality of interaction and collaboration, be it through emotional self-awareness or the heightened awareness of the group members' emotional states [92, 93, 236].

In the context of emotional self-awareness, one of the main examples is the Affect-Aura visualization system [184]. The system uses multimodal sensory information, including audio, video, physiological and contextual logging data, to capture the user emotional state. Note that the sensory information collected this way is obtained from an entire set of both portable and fixed sensors, including a webcam, a Microsoft Kinect device, an EDA, and a GPS device. This information is then encoded visually into a timeline-based representation, where for each hour a glyph is used to encode valence with color, arousal with shape and overall user activity with size.

Similarly, the Affective Diary from [266] promotes affective self-awareness by logging arousal and movement levels through a sensor situated on an armband as well as by recording additional data from the mobile phone (e.g., photographs, text messages, location). Again, the information recorded this way is presented in a visualization that uses mainly colors and shapes to aid reflexion and introspection related to the user's own affective states over the day. In the area of emotion annotation, Ohene-Djan et al. [202] highlight a visualization tool that allows users to express their current affective state in order to reflect their emotional perception about the activities they are currently involved in.

Moving away from the logging context, in [189], contact center employees are encouraged to express their emotions through affective widgets called moodies. The system is based on self-report and allows the user to express their emotional states in the hope of supporting the processing of stressful and negative emotional experiences.

One major domain where emotional self-awareness is gaining rapid ground are games. In this environment, emotions detected through automated methods like facial expressions analysis and BCI systems allow developers to augment both conscious and subconscious user interaction. In conscious emotion-enhanced interaction, users have to focus on obtaining a certain emotional or mental state in order to interact with the game¹. One example for this would be a game that requires you to be engaged or to meditate, be calm, in order to achieve a task. In the subconscious approach, the game estimates the user emotions in order to over visual feedback to the user or execute subtle changes in the gameplay. Further details on the topic can be found in [146].

While emotional self-awareness has been considered as a source of improved interaction and user-centered experience, research has also gone into supporting the social and affective side of collaborative systems through the exploitation of emotional awareness at the group level.

A mobile collaborative system is presented in [235, 236] that focuses on improving awareness of social contexts, stress levels and emotional states. The emotional states and stress levels of the users are self-reported, with the option of augmenting the provided information with sensor data from devices connected to the mobile. The application thus incorporates an instant messaging module, as well as an emotion and context reporting interface including a representation of the group emotional states.

In [274], a solution is presented that combines hardware and software modules for communicating and increasing user awareness of workgroup emotions. The emotional states from the group members are obtained by a combination of self-reports and physiological measurements. The objective measurements relied on the data obtained from probes that were manipulated by the group members, and which reported changes in temperature and vibration. This way, users could compose kinesthetic expressions of emotion, which they could decide to share with the group on public and shared displays. In this particular scenario, group emotional awareness has been linked to conflict resolution as well as enhanced communication and social relations.

¹Games for the Emotiv EPOC headset, <https://www.emotiv.com/store/apps/applications/117/> (February 2014).

A hardware solution is also described in [197], where a digital photographic frame is designed with the purpose of communicating emotions over long distances. The frame, called EmotiPicture, allows users to express their emotional state by offering an online interface where a person situated at a remote location can select an image and a corresponding emotion. Locally, the hardware frame is organized in a similar dual fashion, offering a section for an image to be displayed, and another area for visualizing the corresponding self-reported emotion.

In [252], the Affecter system for emotional reflexion and interpretation is highlighted. The application is based on self-encoding of user emotions through the modulation and distortion of a video feed established between two parties. Following the guidelines of affective interaction more closely, the emotional representations do not follow any rules, users having the freedom of expressing and decoding each other emotions based on their own interpretations.

In [92,93], a set of emotional awareness solutions is presented that are intended for collaborative applications. The range of the proposed solutions spans from GUI widgets that encode the three dimensions of pleasure-arousal-dominance (PAD) for distributed collaboration systems, to affective icons and avatars. Affective icons in the form of a cartoon figure are also employed in [121], where users are asked to give affective ratings to interface elements on websites. All these representations are meant to enhance user awareness of emotional distribution among the group members. In terms of emotion recognition, the affective states are acquired through physiological measurement devices.

Finally, emotional awareness has also been explored in communications systems. In [267], a mobile phone solution is proposed that enables users to augment their text messages with emotional information. The affective states are encoded through the color and shape of the application background, which are mapped to the two-dimensional space of valence and arousal. The user emotions are self-recorded through gesture executed with the device stylus, gestures that have at their core pressure and shaking motion. Related to this, in [227] an email client is highlighted that augments the text of an email with stylized information about the sender's emotional state. In this case, the emotion recognition is powered by facial expression detection and the inspection of user typing speed as an indication for excitement.

When considering emotional awareness as the main purpose of emotion visualization, an important aspect becomes the ability to integrate affective representations in pre-existing interfaces. In terms of augmenting and integrating additional levels of information into these interfaces, researchers have, for example, proposed extensions for the standard user interface (UI) components that could enrich their informational and functional attributes. Baudisch et al. [18] present a graphical approach that briefly highlights those UI components that the user has recently interacted with in order to support improved user awareness without introducing delays or other interference in the interaction process. The topic of awareness is further discussed in the context of collaborative scenarios by Hill et al. [112], where UI components are augmented with additional information about the interactions that other users executed on these widgets.

Besides awareness of interaction, another topic that is closely related to the design of the graphical user interface (GUI) is represented by the distribution of relevant information in highly complex systems or large structured datasets. In such cases, information scents or information cues [56] have the potential for guiding the user to the areas of the system that can offer more relevant information, support his analysis process and improve his decision making. An example in this sense is presented by Willett et al. [313], where GUI widgets are enhanced with additional visual information cues in order to improve, among others, the information foraging process. Note that all the above mentioned approaches aim at offering new information without changing the layout or positioning of the various widgets, thus offering an enhancement and not necessarily an alternative.

While informational augmentation of UI components has its utility, our attention is tuned more towards emotional information. Over the years, different approaches have been proposed to solve the issue of augmenting the interface with emotion visualization solutions. While some are simple enhancements like emoticons [117], others focused on changing the desktop interface by incorporating human-like interfaces like emotional virtual agents [162, 180], adding affective widgets for emotional interaction and relief [114, 189] or designing affective systems [4] that can perceive and react to the user's emotions in application and task-specific manner.

Focusing on BCI detected emotions, Liu et al. [164] employ a commercial mobile EEG headset to detect user emotions and reflect them in the GUI through a 3D virtual agent. Also, Inventado et al. [122] propose a system that supports the easy emotional annotation of user interaction. More closely related to our work, Mehra-bian [185] briefly suggests the representation of emotions in the interface, while Garcia et al. [92,93] analyzes the possibilities for supporting emotional awareness in collaborative systems. One of the solutions proposed by Garcia et al. [92,93], the *emotion awareness graph*, is particularly relevant for seamlessly integrating emotion visualizations in a GUI, as it represents a custom UI component that consists of a histogram-like representation.

In this thesis, we highlight two novel frameworks towards emotion visualization (Sections 4.3 and 4.4). These representations can offer visual feedback about the affective states of the users in real-time and post-task, thus supporting emotional self-awareness and the awareness of other users' emotions. We further show how raising user awareness of social and interpersonal aspects has a positive effect in collaborative scenarios.

2.3 Visualizing the Web and User Online Browsing

While the previous sections of this chapter focused on user-centered aspects involving collaboration, interaction and user subjectivity, in this section we explore the existing visualization approaches for a concrete application: representing online browsing histories.

Maybe the most common approach for visually encoding browsing histories are tree representations. Tools like MosaicG [83], PadPrints [111], Organic Bookmark Management [255], WebMap [71] and Domain Tree Browser [91] use one or multiple vertical or horizontal 2D trees to represent the domain-structure of the navigated websites. In some cases, these tree views are coupled with additional list views that highlight the temporal order of the visits, as the tree representations do not reflect the temporal succession of events. Additionally, in many cases screenshots of the web pages are used as thumbnails embedded in the nodes to support the recognition process [83,111,255].

A comprehensive overview of website log data visualization solutions is given in [183]. Still, all these approaches represent a web page only once in the tree, even if it is visited multiple times, thus mapping the user navigation pattern to the underlying website structure and not the other way around.

An alternative 2D graph representation focuses on capturing and visualizing the branching events in the navigation path [306]. These visualizations manage to capture the sequential aspect of the browsing process, as each accessed page is drawn as an additional node in the graph. If the user navigates back and accesses a different website, the resulting branch will be accordingly represented in the visualization.

A slightly different 2D space-filling solution is offered by the Trails plug-in [316] that supports a hierarchical, chronological and group-based representation of the visited pages. Furthermore, it offers a statistical overview of the most often visited websites. Another method for representing browser histories is highlighted by solutions that employ one [134] or multiple [58] interconnected linear views that are enhanced by graphical elements (e.g., thumbnails).

Besides 1D and 2D solutions, web browser histories that employ multiple dimensions or intuitive metaphors have been developed. VISVIP [62] is a 3D representation of a navigation log, where two dimensions are used for drawing the website structure, while the third one encodes the temporal information. On the other hand, the combo WebBook and WebForager [40] use the concept of a book to give an overview of the websites as well as offer an intuitive information-space for the user.

A special class of browser histories is represented by the statistical summary histories. Tools like SlifeWeb², RescueTime³ or Eyebrowse [292] are mainly focused on time management and analytics, and allow users to generate their own statistic view about how they—or others—navigate the Internet. Focused on improving website usability, Eick [77] developed a suite of interactive visualization solutions for exploration, analysis and pattern recognition on website activity logs. In parallel, users now have the option of employing web browser extensions that focus on online browsing history visualization. One such example is the Yogurt extension [301] that uses a barcode

²Slife labs time management software, <http://www.slifeweb.com> (February 2014).

³Rescuetime, <http://www.rescuetime.com> (February 2014).

representation to visualize the access history to a set of predefined websites over a time period of multiple days. A similar yet simpler solution is offered by the Iconic History [120] browser extension that visualizes the website access history of a user by generating a continuous favicon stack.

However, browser histories are not the only type of data revolving around complex, interconnected temporal events. Other time-series visualizations employing similar visual concepts to our approach include World Lines [305], a visualization technique for exploring the alternative paths of heterogeneous simulation runs, as well as Cloud-Lines [153] and LeadLine [302], time-series visualizations for identifying and representing meaningful events, e.g., in news and social media data.

While diverse and functional, none of these methods focuses on the complex parallel browsing habits of today, where tabs and windows have become means for the user to organize thoughts, actions and accessed information [119]. The importance of a tool for visualizing, analyzing and comparing parallel browser behavior is further highlighted in [10], since users tend to use multiple windows and tabs as means for backtracking (e.g., users abandon the use of in-browser back operations in favor of opening new tabs and switching between them) and multitasking (e.g., users interact with one tab while web pages are being loaded and processed in others). Similarly, findings from [265,307] suggest that users often employ parallel browsing in web search tasks for reasons like comparing search results, executing multiple queries, interacting with a page while others are being loaded, etc.

In the following chapter, we introduce an interactive visualization tool that aims to address and highlight user parallel browsing behavior in online browsing histories. Moreover, the resulting time-series visualization of semi-hierarchical data will be used further to test and evaluate our user-centered collaborative visualization techniques, proposed in the Chapters 4 and 5.

Chapter 3

Visualizing Parallel Web Browsing Behavior

Today, web browsers are the main software systems that are used for accessing the information and functionality of the Internet. As a result, web browsers have grown in complexity, in order to support a wide range of applications, including many from the fields of communication, games, entertainment, development, etc. Given this variety and complexity of the tasks that can be executed online, web browsers implemented a set of parallelizing features that allow users to visit web pages in multiple concurrent threads. The most relevant of these features is the support for multiple browser windows and tabs. As a result, users can now execute multiple online tasks at the same time as well as switch back and forth between these web pages. This in turn is called *parallel browsing behavior* or *tabbed browsing behavior*. However, while the functionality that supports the parallel exploration of online websites is already being exploited, there are only a limited number of solutions that allow users to detect and analyze browsing patterns in parallel browser histories.

In this chapter we introduce *WebComets*, an interactive visualization system for exploring multi-session multi-user parallel browsing histories. WebComets is aimed at researchers, usability practitioners, and other professionals for the analysis of larger parallel browsing datasets and the extraction of relevant online user activity. Furthermore, the size and complexity of user parallel navigation logs, especially when inspecting data for multiple sessions and users over a long period of time, makes such a visualization system the ideal platform for testing and evaluating our UCCV techniques, detailed in Chapters 4 and 5.

The remainder of this chapter is organized as follows: First, we highlight the motivation behind this research. This is followed by a requirement analysis, a detailed discussion of the design decisions and the supported interactions of our proposed visualization. We then introduce a motif-based contextual search for enabling the filtering and comparison of user navigation patterns. In order to validate our approach, we describe an evaluation of our tool, and lastly we offer our conclusions. Our investigations suggest that parallel browser history visualization can offer better insight into user tabbed browsing behavior and support the recognition of online navigation patterns.

The following sections are based on the publications entitled “WebComets: A Tab-Oriented Approach for Browser History Visualization” [50] and “An Interactive Visualization for Tabbed Browsing Behaviour Analysis” [51]. Additional details about this research project can be found in the bachelor’s thesis “Interaktive Visualisierung erweiterter Browserchroniken” [285] by Igor Truderung and the master’s thesis “HistoryLane: Web Browser History Visualization Method” [57] by Igor Chtivelband.

3.1 Motivation

Since the release of the World Wide Web to the public, the Internet instantly became an important source of information as well as a communication platform without which today’s world is hard to imagine. However in the last decades, the tasks that users could execute on the web have also greatly increased in complexity, thus influencing the range of features that today’s web browsers incorporate. As a result, users now employ the browser in extremely diverse scenarios, ranging from checking e-mails, chatting and streaming media, to playing games, managing schedules and even developing software [292].




Figure 3.1: Multiple tabs open in the same web browser window, allowing the user to access multiple websites and execute multiple tasks in parallel.

To meet this tendency, many web browser developers started implementing *tabs* to allow users to access and explore multiple web pages simultaneously inside the same browser window (see Figure 3.1). While similar, this experience is still significantly different from the one of using multiple browser windows, as the more lightweight tabs allow users to have an overview of the opened web pages and even organize (e.g., group) the various tabs by the topic of the loaded page.

The importance of tab-based operations—also called online parallel browsing behavior—can be further supported by the work of Miyata et al. [190], where the presence of foreground and background tasks and their interconnection in the user’s mind are emphasized from the perspective of cognitive psychology. This has been also recognized in 2011, when the W3C specification of Page Visibility API was proposed that would allow developers to determine which web pages are in the foreground from a set of tabs or windows. As such, web browser tabs are specifically designed to follow this principle and allow users to distribute their attention based on this model. Furthermore, over 97% of today’s web browsers market share support online parallel browsing in the form of windows and tabs¹.

Recent studies on the topic of tabbed browsing have also captured a clear tendency. Labaj et al. [155] inspected the potential of improving user models by analyzing parallel browsing logs and exploring user habits. In the work of Dubroy et al. [73], a strong user bias towards using tabs instead of multiple windows is highlighted. At the same time, tabs are described as a technique that supports task management, with the following categorization of common tab usage patterns: reminders and short-term bookmarks, opening links in the background, multitasking, and comparing. Huang et al. [119] highlight that 57.4% of Internet sessions in the browser make use of tab-based parallel navigation. Other findings comprise that these values fluctuate between 4-85% [294]. While not conclusive, such a wide range suggests an even more important aspect: there is currently only limited information and insight into the way users explore and organize their parallel browsing experiences, and a visualization tool would be needed that is capable of capturing and reflecting the intricacies of online navigation today. Such a powerful representation, enabled by meta-information about user online sessions, has the potential to simplify the tasks of researchers and analysts in fields like information retrieval and behavioral sciences.

¹W3Schools’ browser statistics, http://www.w3schools.com/browsers/browsers_stats.asp (February 2014).

Sadly, in most cases the representation of the history data is limited to a textual list of website names and URLs that is sorted based on a certain criterion, e.g., chronological or by load frequency. Additionally, these browser histories give little insight in the call hierarchy of the web pages, the relevance of a particular site to the users or the back-forward operations they executed [105]. In other words, while browsers offer support for parallel browsing, most current histories fail to capture this aspect [119] and thus do not reveal any interconnection patterns between the web pages or user sessions. Thus, as temporal features are not sufficiently highlighted in conventional histories, one cannot recognize the connections between websites and browser windows and tabs. This information is relevant in a setting where users now engage in parallel navigation and distributed attention processes between many open browser tabs and windows.

In the following, we address the problem of designing the WebComets interactive visualization for Internet browser histories that allows analysts to quickly find, compare and analyze parallel navigation behavior based on a set of existing—e.g., those described in [119]—and novel metrics as well as supports intuitive search operations based on content and context information.

3.2 Requirement Analysis

To support flexible search and analysis efforts, all control and interaction elements that users employ to organize complex and parallel browsing sessions—such as windows, tabs and back-forward operations—need to be recorded and graphically represented in a first step. The captured data organized into a set of different user profiles will be at the core of the later visualization. It has to embed rich meta-information that could be of interest for the user of WebComets, i.e., for researchers in information retrieval, behavioral sciences and related fields. Following the study described in [119], typical research questions for a better understanding of parallel browsing behavior on the web are for example: *When and how often are users parallel browsing on the Internet? What affects parallel browsing behavior during interaction with web search results?* Another possible research question is to identify reasons why users revisit pages—for example, because of monitoring pages [1, 139]—and how this is typically done in a multi-tab browser environment.

For supporting such studies and for finding answers to such questions, a visualization tool has to offer specific functionalities. We have composed a list of requirements that need to be satisfied by a browser history visualization based on preliminary studies, user feedback about browser histories and their limitations, and information about the nature of parallel browsing behavior from the previously referenced publications. Fundamental requirements for the visualization of the captured data are listed in the following:

1. The visualization should offer an *overview* of the loaded data and *support detailed investigations*. This can be achieved by means of tailored interactions and methods like “details-on-demand” [258].
2. The *temporal flow* of the visualized navigation sessions needs to be clearly distinguishable. While most web browsers only display a chronologically sorted list of the accessed web pages, it is important that the temporal dependencies are visually and relationally highlighted.
3. For each visited web page, *additional meta-data* needs to be captured like the duration of each visit or the duration of interacting with the web page. While many solutions already count the number of executed accesses to each web page, it is important to visualize the temporal sequence of events related to any sequence of navigated pages because documents may be inspected more often and for longer periods.
4. *Website categories*, like search engines or news websites, should be introduced and represented. Current browsers support bookmarking of web pages, which implies saving the web page’s address while at the same time tagging it with the help of keywords or categories. These operations increase the retrievability of stored bookmarks and should also be offered by our system.
5. The visualization should clearly represent which *navigation path*—or sequence of visited web pages—the users have followed during their browsing sessions. This includes information about what browser windows and tabs have been opened and closed. Thus, it will be possible to reconstruct the steps that lead to the opening of a particular web page—a missing feature in many related solutions.

Requirements that improve scalability (with respect to log size and number of user profiles) and analysis possibilities:

1. *Visualizing multiple browsing histories* at the same time should be supported in order to allow comparison and analysis operations (e.g., detect if multiple users have similar interests).
2. *Connections between similar websites* should be emphasized as these might be relevant alternatives in search and analysis tasks.
3. The analysts should be able to *search for particular web pages* based on content (e.g., title or category) and context information (e.g., pages accessed prior to the one in question). While content-based search is present in all history lists, a search for the context is not supported in most cases.
4. Equally important for the analysis of parallel browsing behavior is the *search for navigation patterns*, i.e., finding specific structures (motifs) in the navigation graph which results from branching out from a linear navigation behavior by using tabs or additional browser windows.

3.3 WebComets: Tab-oriented Browser History Visualization

WebComets is a system for the interactive visualization of extended, tabbed browser histories. It was implemented in Adobe Flash ActionScript due to its rapid prototyping abilities and the flexibility it offers in terms of online and standalone capabilities. This was important in an initial step of the research, where we planned to get feedback about our design also from people situated at remote locations. The representation and interaction metaphors of WebComets satisfy the requirements highlighted in the Section 3.2. Figure 3.2 shows a screenshot of our tool.

To achieve the required functionality and analysis capabilities, WebComets cannot solely rely on information gathered by standard logging systems. For example, browsers like Mozilla Firefox or Internet Explorer do not record the duration for which a user has actively interacted with a web page. More importantly, browsers do not focus on capturing the parent-child relationships between accessed web pages and even less the connections between opened tabs or windows. Other researchers have




Figure 3.2: WebComets visualization of the parallel browsing histories of two users (light-brown and light-grey background). Each horizontal line represents the timeline of a tab that the user has opened, while vertical branches highlight new tabs that have been created by clicking a hyperlink in the parent tab. The comet-shaped glyphs encode loaded websites and their color coding represents topics. Their position on the time axis depends on the moment when they were accessed. In the current visualization, one can notice that the first user had up to six tabs open within the shown timeframe. The second user employed two browser windows without additional tabs and spent most of the time on one web page as represented by the long comet tail in the second window.

also encountered this difficulty when investigating the parallel browsing behavior of Internet users [119]. The unavailable information included, among others, missing source tabs for branching operations and no information on how a tab or window was created (i.e., new tab or hyperlink click).

To address this, we developed a Mozilla Firefox browser add-on with the help of Javascript and libraries like jQuery² and KineticJS³. The add-on incorporates the ability of recording and saving an Extended Browser History (EBH) inside an SQLite database. The information collected in this manner can be subsequently accessed and visualized for multiple browsing sessions and users.

²<http://www.jquery.com>, jQuery JavaScript library (February 2014).

³<http://www.kineticjs.com>, KineticJS JavaScript framework (February 2014).

For any current user profile, the extension saves the navigated URLs together with relevant additional information. The complete EBH includes a subset of the metrics employed in [119] and a set of additional metrics relevant to the analysis and comparison of parallel browsing habits and behavior. Thus, for each user profile, the EBH records the following information:

- user profile information (such as username),
- opening and closing times for each tab and window (tab and window sessions), as well as
- type of creation for each tab and window, i.e., opened blank or through a link from another tab or window (branching).

To complement this, the following data will be recorded for every accessed web page:

- the title of the web page including its URL,
- information about how a web page was opened (i.e., through direct input of the URL in the address bar by the user, through clicking of a link in another web page, or through the browser's integrated back/forward operations),
- category of the web page based on a static list of web domains,
- number of accesses to a domain and a particular URL (pageview),
- additional time intervals (i.e., total time representing the time interval when a web page was loaded and discarded; focus time representing the time interval for which a web page was in the foreground; and active time representing the time interval for which the user interacted with a web page).

Note that the tab switches metric proposed by Huang et al. [119] is currently not being stored in the EBH, as we argue that additional time intervals offer an alternative view for the distribution of the user's attention over multiple tabs and windows. In the following, we highlight the visual design and interaction capabilities of WebComets together with supplementary information about the EBH.

3.3.1 Visual Representation

In order to satisfy the requirements highlighted in Section 3.2, the WebComets visualization has to consider a variety of aspects, including the design space analysis for graphical history tools presented in [105]. Probably the most important one is the representation of the temporal dimension and the mapping of the web pages to a time axis. In order to use the larger width to height ratio of modern screens (widescreens), a visualization concept was devised that maps the timeline to the horizontal axis, from left to right. In this representation, each accessed web page is displayed as a circle and gets assigned its corresponding position on the timeline (x-axis), see Figure 3.3.



Figure 3.3: Conceptual representation of a tab timeline that maps the flow of time from the left to the right. The filled circles on the line represent visited web pages. Their diameter may represent different attributes.




Figure 3.4: Conceptual representation of the tab hierarchy. Each horizontal line represents a single tab. Tabs that are connected through a vertical line suggest a parent-child relationship between the tabs and, respectively, between the web pages loaded in the tabs at that point in time.

Another important aspect of the visualization is represented by the encoding of the parallel navigation of the users in multiple browser windows and tabs. WebComets represents each browser tab as a separate horizontal line segment that is parallel to the time axis (Figure 3.4). This combination of patches and parallel segments is similar to the representation of a parallel browsing session in [119], as well as to [153] where multiple time-series are visualized through a comparable solution.

As tabs can be opened manually or by clicking a link in another tab, this can result in a tree-like structure that suggests connections in terms of hyperlinks, but also potential themes between various websites. This parent-child relationship is represented in the visualization as two horizontal lines connected by a vertical one (Figure 3.4). At the same time, if the user opens a tab manually, there is no clear way of connecting the first web page of this tab to any other already opened pages. Therefore, a new tab line is shown as disconnected from the rest of the tabs that were already loaded.

At the same time, multiple opened browser windows are visually encoded as framed rectangular areas, where each rectangle contains a tree-like structure of tabs that reflects the opened tabs in each window during the user session (see Figure 3.2). As a rectangle stretches along the horizontal axis, its left and right margins represent the opening and closing times of the window. Note that a rectangular representation can be also activated at the tab-level to enforce the navigation patterns of the users. All rectangular shapes have a specific background color that identifies them as windows/tabs belonging to a certain user profile.




Figure 3.5: Circular glyph representation of a visited web page. The figure highlights two versions of representing the focus (dark blue) and active times (light blue) of the visited web pages: as a comet tail (left) or as beams (right). The analysts can switch between the two representations, as the comet tail emphasizes the web pages where the users spent a lot of time, while the beams allow for a more scalable representation.

The representation of a visited web page is encoded as a visual metaphor [140], where each page is represented through a comet-like glyph. Each such glyph has at its core a pie chart enriched with additional graphical elements. These elements encode multiple EBH attributes, as shown in Figure 3.5. The circular glyphs are mapped on the horizontal axis to the moment in time when the corresponding web page was loaded, while the vertical positioning identifies the host tab where the web page was loaded. The radius of each circular representation encodes the number of visits (c_{site})

the user executed in the current session to a particular domain, e.g., *www.google.com*. To correctly visualize this, the maximum visit counts (c_{max}) are computed for each domain in the EBH. Also, users have the option of setting and storing the minimum and maximum size for the pie chart radii (r_{min} and r_{max}). Based on these values, the radius of each web page circle is computed as:

$$r = \frac{(r_{max} - r_{min})(c_{site} - c_{max})}{c_{max} - 1} + r_{max} \quad (3.1)$$

As a result, pages that are part of domains that are visited more often will have a larger pie chart than those that have domains that are visited rarely.

At the same time, the pie chart representation is divided into two sectors. The first sector to the right captures the ratio of visit counts for the current web page compared to the overall visit count for the domain. For example, if the domain *google.com* has been accessed six times in total and the current web page (*www.google.com/search?q=conference*) only two times, then the pie chart will encode a sector of 1/3. This is computed by the following formula where c_{link} represents the number of accesses executed to the current link:

$$\alpha = \frac{360 \cdot c_{link}}{c_{site}} \quad (3.2)$$

As mentioned previously, the EBH includes three time intervals for every web page: total, focus and active time. Note that the focus time stores the amount of time the web page was in the foreground and active time captures the duration for which there the user was clearly present and interacting with the web page, e.g., by mouse movements or key strokes. Therefore, the following relationship is valid for any visited web page:

$$total_time \geq focus_time \geq active_time \quad (3.3)$$

The total time for a web page is represented by the horizontal segment between its pie chart representation and the position of the pie chart for the following web page on the same tab. Because of Eq. 3.3, the focus and active times are visualized as subsegments of the total time. This is achieved by two representations that resemble a comet tail and a beam that follow the pie chart on the right side (Figure 3.5). Both intervals have their origin at the loading time of the current web page. The analysts

can switch manually between these two encodings of the focus and active times, as our studies have shown that both representations have their advantages: the comet tails are better at emphasizing the web pages where the users spent a lot of time, while the beams have a more concise representation that is more scalable (see Section 3.3.4).

The length along the timeline is computed by Eq. 3.4, where t_{enter} and t_{quit} are the timestamps for the begin and end of the session; $x_{rightEdge}$ is the rightmost x position and Δt is the focus time and active time, respectively.

$$l = \frac{x_{rightEdge} \cdot \Delta t}{t_{quit} - t_{enter}} \tag{3.4}$$

















SAMPLE	CATEGORY	SAMPLE	CATEGORY
MULTIMEDIA		INFORMATION	
	MUSIC		NEWS
	RADIO		SPORT
	VIDEO	PRODUCTIVITY	
	LIVE STREAM		SEARCH ENGINES
	IMAGES		SOFTWARE
COMMUNICATION			BANKING
	CHAT		SHOPS
	SOCIAL NETWORKS	OTHER	
	E-MAIL		GAMES
			UNKNOWN

Figure 3.6: The list of supported categories together with the corresponding colors for the circular encodings.

The glyphs are colored based on a correspondence table to a preset number of website categories (Figure 3.6). The categories support the quick differentiation and detection of various online user activities. They have been generated based on [101,138], where common in-browser activities are highlighted. For each visited link, a category is

selected by the Firefox extension by inspecting the domain of the current website. The process employs a database of domains that is divided into 15 categories. The lists of relevant domain names for each category have been generated with the help of the Alexa Top 500 websites database⁴. If a domain is not found in the database, the web page will be included in the “unknown” category (see Figure 3.7). In order to make the different categories as distinguishable as possible, the color selection is based on recommendations from the Colorbrewer website⁵ [30]. Moreover, the analysts can inspect a legend located in the visualization space on the left hand side. By clicking on a category in the legend view, it is possible to select all web pages from that group.



Figure 3.7: Example of a web page comet from the category *unknown*.

A common approach for bookmarking systems in browsers is to store not only the URL and page title, but also the favicon of a web page. A favicon is a small (commonly 16x16 pixels), squared icon that identifies the domain of the web page and usually appears next to the address bar of a browser when a web page is loaded. To improve the chances of a web page being quickly recognized [135] in the WebComets visualization, the analysts have the option to additionally display favicons over the pie charts, if these icons are available (Figure 3.7). Due to the small size of the icons, the sectors of the pie chart remain distinguishable. The size of the favicon is proportional to the one of the circle. This means that for small circles, the icons are also reduced in size. To compensate, users can hover over the icon with the mouse cursor in order to represent it in its original resolution.

⁴Alexa: The top ranked sites in each category, <http://www.alexa.com/topsites/-category> (February 2014).

⁵Colorbrewer: Color advice for maps, <http://www.colorbrewer2.org> (February 2014).

When navigating the Internet, users have multiple options to reach a certain URL: type it in themselves, click on a hyperlink, execute browser-specific operations, etc. While some operations suggest relationships between web pages (e.g., hyperlink connections), others might represent a breakpoint in the thought process of the user. Such a case is usually given when users type in a new URL themselves. To better highlight this, web pages that are loaded in such a way have a short vertical bar sticking out from under their glyph to suggest a possible mental breakpoint.

Users may access a web page by utilizing the back and forward navigation buttons of the web browser. In many modern browsers, these buttons allow the user to navigate backward and forward one or multiple pages. Such operations are useful especially in cases where the users feel they have reached a “dead end”. If a web page was loaded through a back or forward step, this is captured and represented in the visualization through small arrow-shaped elements on the left or right of the corresponding pie chart (see Figure 3.5). The numbers #b and #f inside the arrows highlight how many back or forward operations were executed at once in order to reach this web page.

3.3.2 Supported Interaction

A set of interaction metaphors enhance the visualization’s abilities by implementing Shneiderman’s information seeking mantra [258] and by addressing issues like flexibility and scalability. Besides giving an initial overview of the loaded browser logs, the WebComets interface supports pan and zoom operations, similar to modern interactive maps. While the panning operation is self-evident, there are two zooming approaches implemented in the tool: one is a regular 2D zoom that allows analysts to inspect the details in a certain area, while the other is a 1D horizontal zoom along the timeline that stretches the horizontal axis (Figure 3.8).

When a web page is selected, WebComets searches the browser history to check for similar web pages. By default, only pages with the same domain are considered to be similar, but more complicated rules for interconnection can be generated, e.g., pages containing a keyword, pages that have a similar active time value, etc. These are highlighted by adding links between each selected element and its counterpart. The curved lines (Figure 3.9) are used for showing the presence of similar web pages




Figure 3.8: Temporal zoom (1D) along the horizontal time axis: original zooming factor (left) and 30x horizontal zoom (right).

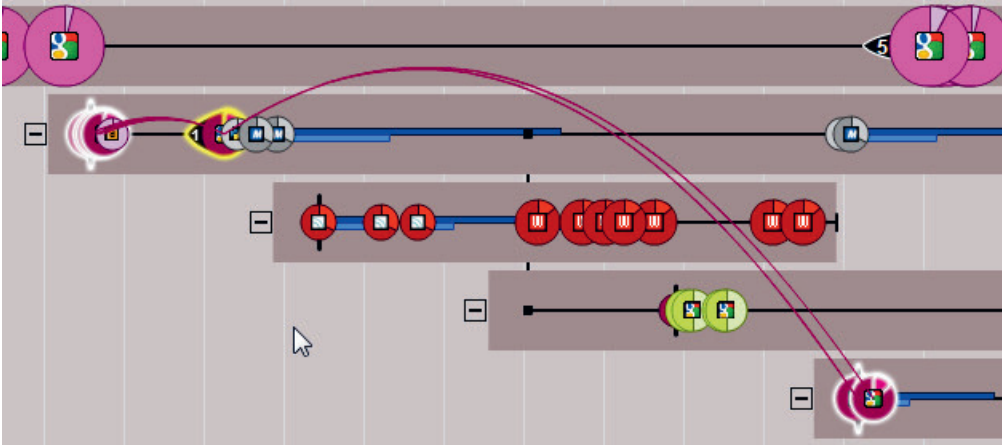


Figure 3.9: Connections between the selected web pages and other glyphs highlighted through continuous curved lines.

to the selected ones, possibly in areas of the visualization that are currently not visible. Curves can be easily perceived as they contrast with the overall orthogonal representation of our approach. To avoid clutter, curves are drawn in such a way that the probabilities of intersecting lines or a curve intersecting a glyph are reduced. More precisely, our tool computes a predefined set of possible curves with different curvatures to connect the two nodes and then utilizes a collision detection algorithm to identify the curve that intersects the least number of glyphs. This is achieved by drawing up to 20 consecutive Bézier curves for each pair of nodes. Initially, a curve is computed with default control point values. If this curve intersects the circle of other nodes, the control parameters are modified along the vertical axis in order to increase the curvature of the connection. For the new curve, again we inspect the overlap with

other nodes. This process continues either until a curve is found that does not overlap other nodes, or until 20 curves have been computed. In the later case, the algorithm simply selects the curve parameters for the curve that had the fewest collisions with other nodes.

It might occur that multiple glyphs are partially or almost totally overlapping. Even if the analysts have the possibility to execute timeline zoom commands to compensate for this and clearly separate the overlapping glyphs, this is a vital scalability issue. To deal with this, WebComets displays partially overlapping pie charts by positioning the glyph of the web page that has been accessed later on the top of the previous one. To further compensate, the analysts can move the mouse pointer over a set of densely grouped circles. By doing so, the glyph with the center closest to the pointer will be moved to the foreground. This focused glyph is additionally complemented with graphical elements (e.g., favicon, if not enabled globally) and textual information (e.g., web page title). In cases where the analysts wish to focus their attention on the temporal constraints and relationships, all glyphs can be reduced to dots, minimizing the overlap of glyphs and tails. Furthermore, WebComets detects instances where two or more glyphs are completely overlapping and replaces these with the one that has the largest radius. In order to suggest the presence of other circles underneath it, the border thickness of this circle is increased proportionally with the number of hidden circles. Besides controlling the depth order of the elements, moving the mouse pointer over a glyph opens a tooltip that shows the title of the web page and its link as well as a set of customizable page attributes (e.g., title, URL, or pageviews).

Selecting a glyph can be done by mouse click, or in order to select multiple glyphs, by holding the Shift-key while clicking. Besides this, analysts can also use the legend to select all elements from a category or search for certain patterns to highlight glyphs that satisfy the search rule. Selected web pages are highlighted and an information box opens in the upper-right corner of the visualization. The information box contains the attributes of the selected web page, shown in Figure 3.10. This box can also be customized in order to include all attributes or a subset of them. In cases where multiple web pages are selected, the information box displays initially only a list with the titles and URLs of the selected pages. If an analyst moves the mouse cursor over

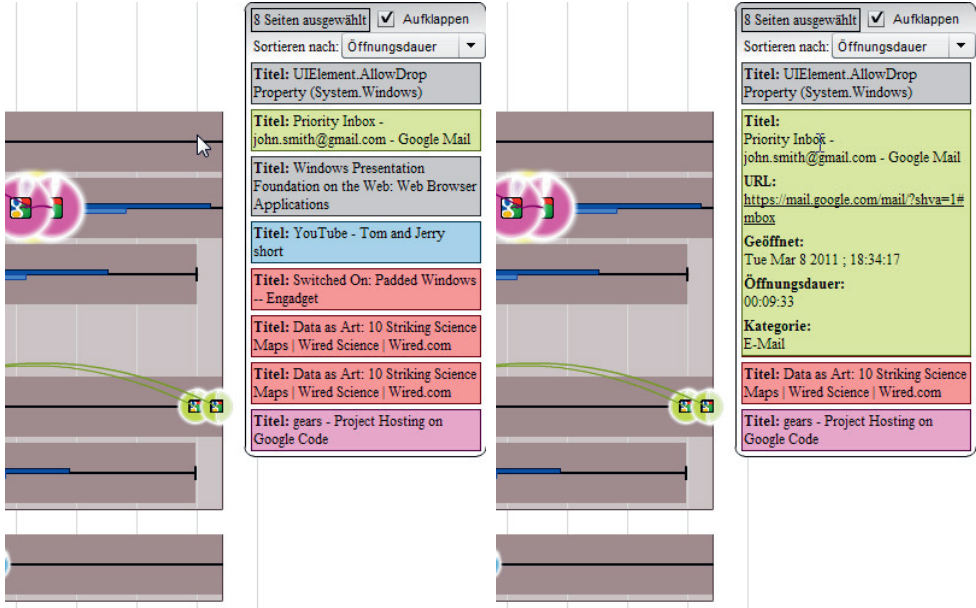


Figure 3.10: Information box presenting details about the selected web pages. By moving the cursor over the list elements, additional information is displayed for the current web page (right, in light-green).

the list, the selected list element will be expanded to present all available information of that page. The background color of each list element matches the category of the web page. A click on the link will open the corresponding web page in the user's default browser.

The analysts could find some parts of the visualized history more interesting than others. Due to this, but also in order to save screen space and offer a better overview, they have the possibility to collapse tabs or windows that are not relevant in the current session. For collapsing the representation of a tab, the analysts have to click the plus icon next to its spawning point; the same is valid for windows. In cases when a tab is collapsed that also has other tabs created by it, all branches of that tab will be compacted together with the parent tab, and vice versa for expansion. Nonetheless, if a glyph is selected and similar web pages were detected on any collapsed tab or window, then the corresponding elements are still visible by links and additionally small dots on the tab lines (see Figure 3.11).

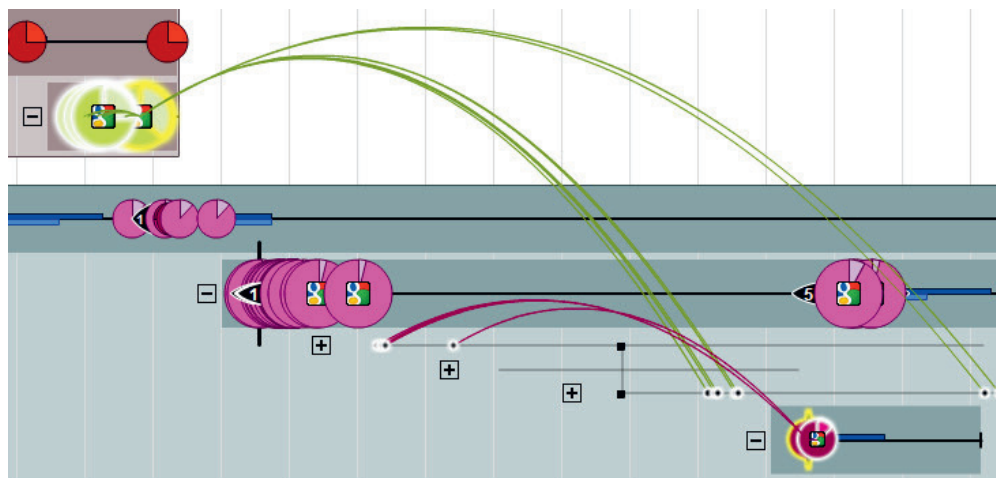


Figure 3.11: Collapsed view. Query-related web pages are also highlighted on the collapsed tab lines.

Besides the already presented features, analysts have additional option to customize the level-of-detail in the representation. Certain features—like favicons and back-forward arrows—are initially disabled and can be included on demand. This allows analysts to maintain an overall reduced level of complexity for the glyphs and the structural data, and only access certain information when these become relevant. The current configuration for the level-of-detail can be stored inside an option window and includes—but is not limited to—the following: switching between comet tail and beam representations, customizing min and max values for the pie chart radii, enabling or disabling web page attributes like favicons or back-forward arrows, and selecting the EBH attributes to be displayed.

3.3.3 Content vs. Context: A Motif-based Filtering

Highlighting different elements in a browser history is closely coupled with searching for web pages or navigation patterns. Most web browsers support a text-based search of their records that limits their ability to detect context information. This limitation is however overcome by WebComets. Users can search for terms and keyword combinations (e.g., *apple+pc*), strict phrases (by using quotes, e.g., *"apple pc"*), or even exclude words from their query (by using the minus sign, e.g., *apple-pc*).

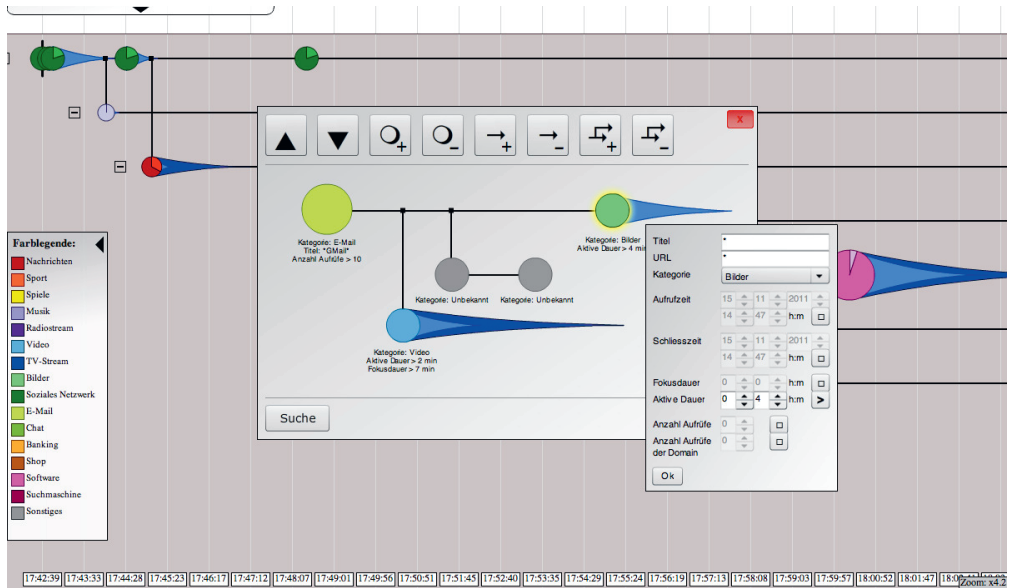


Figure 3.12: The motif search window helps analysts construct, save and search for custom navigation patterns based on web page attributes and context information.

While supporting a text-based search of the extended browser history is vital, the detection, analysis and comparison of temporal patterns in the navigation graph requires an entirely different approach. Thus, WebComets offers a motif search window (see Figure 3.12) that allows analysts to define, store and search for custom information and patterns of navigation. Inspired by the building blocks concept in [224,296], motifs in WebComets are custom-defined subgraph structures that can be used for filtering the current history based on context. As suggested by the buttons in Figure 3.12, analysts can not only create custom motifs, but also store and load them at a later point in time.

These subgraphs can be generated in two ways: either by mining substructures from the currently opened history log files or by manually defining a motif and its corresponding rules. In the first case, the analysts would look through the visualized browser histories and select any subset of glyphs that could be at the core of a parallel browsing pattern. When all relevant elements are selected, the generated motif can

be edited further in the motif window to generalize or particularize the final structural pattern. Contrary to this, in the second approach the analysts would start generating a motif by directly building it in the motif window, adding node after node and customizing their attributes based on their experience and assumptions.

When creating a custom motif, users can add information about two aspects of the subgraph: structural information and node (i.e., website) attributes. For the structural information, users are limited to the navigation patterns commonly supported by tab-enabled web browsers. This means that, guided by the EBH information, users can create individual tabs (connections) by adding a horizontal line to a node, or dependent tabs (parent-child connections) by adding a vertical line that connects a node to another tab line (see Figure 3.12). On the other hand, double clicking on any node opens an additional window that allows the user to add or modify all the attributes of the node/website considered by the EBH database. Note that for numerical attributes, the option exists to define not only exact values but also intervals and restrictions (e.g., active time more than 4 minutes).

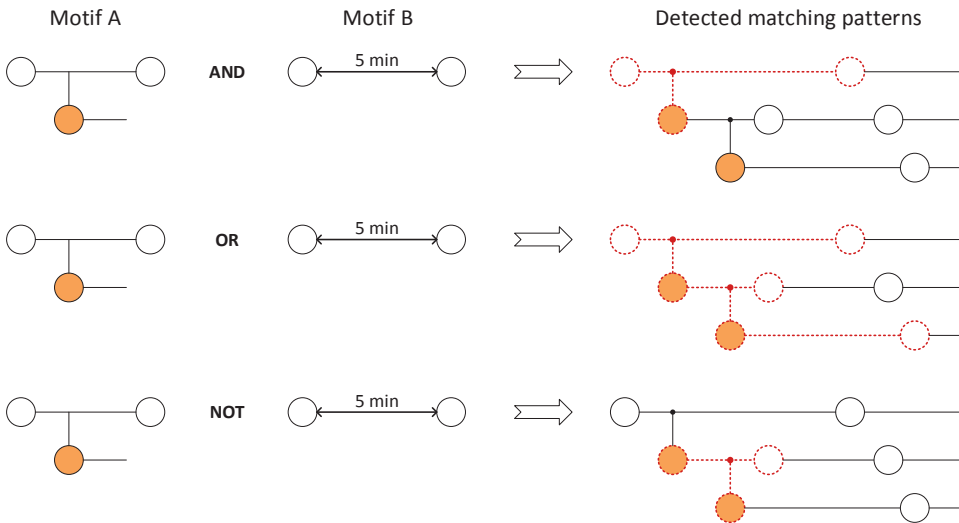


Figure 3.13: Example of logical combinations of custom motifs (left) and their corresponding search results in a parallel browsing log (right).

When compared to other approaches [170], the WebComets motif search has the advantage of allowing logical combinations of sought patterns (e.g., find all node groups that satisfy motif X and do not satisfy motif Y). Figure 3.13 offers a concrete example in this sense. Note that in this figure, motif A encodes two restrictions: a structural one represented by the branching of a new tab, and a node attribute illustrated by the orange color of the first website on the new tab, suggesting that this website should be in category “banking”. Motif B encodes only one restriction, that the total time attribute of the node should be at least 5 minutes long. These two motifs are then combined with logical operators and their resulting motif is used as the base for a complex query. The results of the query on a sample parallel browsing history are highlighted on the right side of Figure 3.13.

As such, users analyzing an EBH can not only look for the topics of the visited pages, but also detect intricate navigation patterns. For example, the motif search could detect that in 72% of the cases when Internet users accessed their e-mail accounts, they also opened a web page from the category “video” in a new tab originating from their e-mail page. This might suggest that these users received e-mails with links to video content. Thus, filtering the browser history based on structural aspects can have many applications, as for example, detecting a web page where the analysts know some attributes of the originating site, or investigating similar interest and patterns of navigation between multiple users.

It is also possible to filter out numerical and temporal values by giving exact numbers, suggesting min or max thresholds, or defining intervals. These rules are then incorporated in the motif and displayed under the corresponding pages. In terms of structure, complex motifs can be built by adding multiple web pages and highlighting existing relationships between them, be it on the same tab/window or on different navigation branches. Once the motif specification is finished, the analysts execute the search operation and the sum of all the rules will be used for filtering the history. Finally, the nodes that fit the query will be highlighted as already described.

3.3.4 User Study

Following the guidelines highlighted in [41], an evaluation of the WebComets visualization has been executed in order to capture its advantages and limitations. The aim of the study was to compare the participants' performance and accuracy when inspecting and comparing patterns in multiple parallel browsing histories. For this purpose, the participants would interact with the same EBH log files by two different approaches: the WebComets visualization and a list-based browser history. The evaluation involved 20 participants with experience in knowledge exploration as well as extensive background in using diverse web browsers and accessing a variety of online applications. All participants had prior knowledge and at least some experience with list-based browser histories.

As an initial step, the participants were randomly divided into two groups of equal sizes. Each member of the first group would have to solve a task using the WebComets visualization, while the members of the second group used a list-based history for the same task. Next, the functionality of the two tools was highlighted to each member of the corresponding groups. Note that the data contained in the browsing histories were almost identical in content, except for the fact that the list-based history was not able to represent the additional fields generated and included in the EBH.

The scenario involved the analysis of two browsing sessions from different users (see Figure 3.14). An initial assumption was made that the users have participated in an online conference call where they suggested relevant websites to each other. This collaborative browsing approach is frequently used in cases where one party tries to highlight some information to the other. A simple example for this would be the collaboration between two students that are preparing for an exam at remote locations. As a result, both users would access multiple similar or identical web pages in the same time interval. The test persons had to determine if the initial assumption of collaboration is supported by the browser histories, and if so, what web pages might have been involved and in what time interval this collaboration took place.

The results were evaluated by inspecting the total time participants took to find a solution and the time frame they reported as part of the conference session. On average, participants managed to find a solution more than twice as fast with the WebComets visualization than with the text-based representation of the navigation

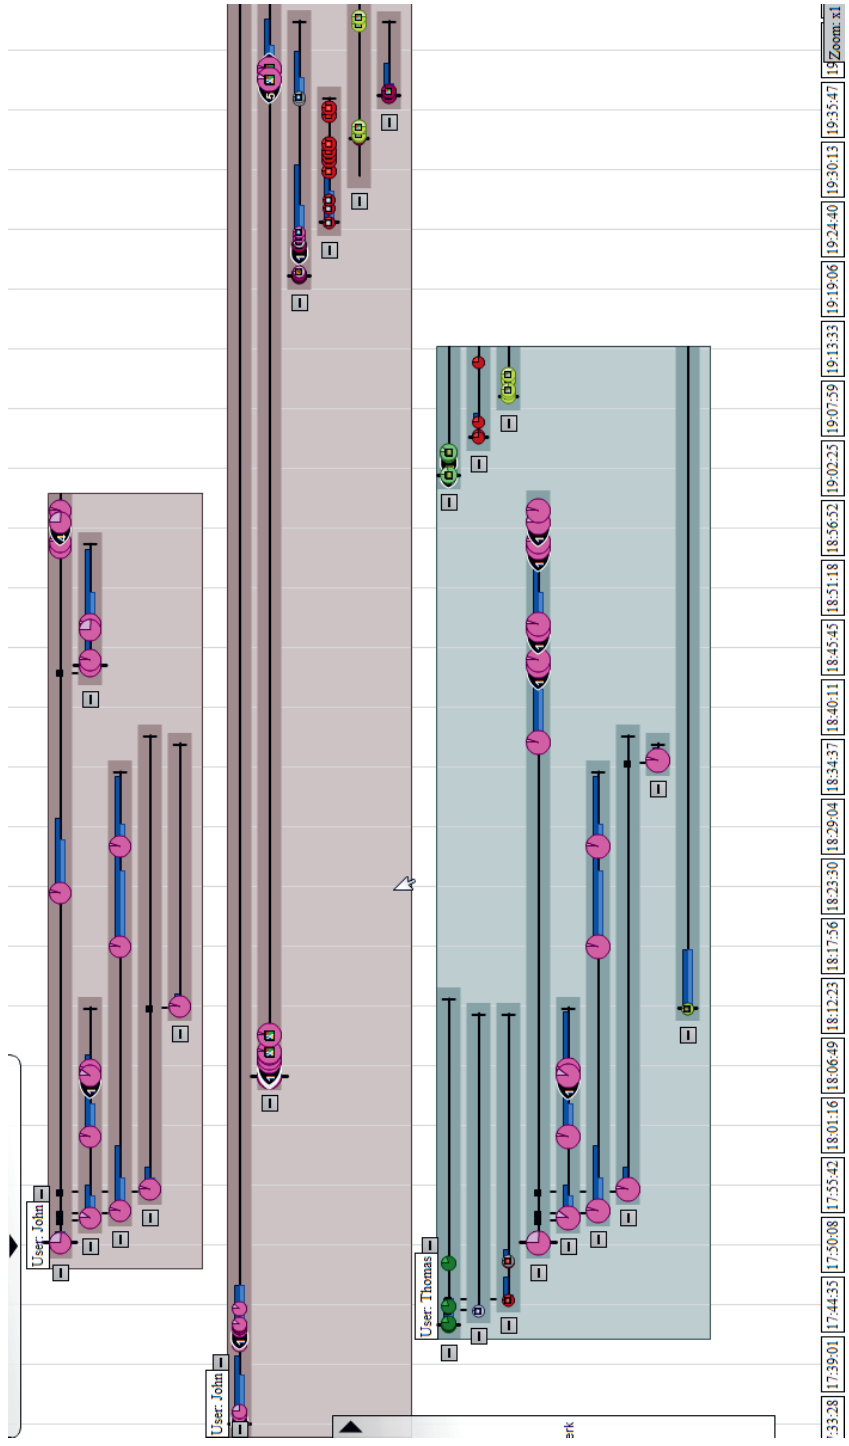


Figure 3.14: WebComets visualization of multiple parallel browsing sessions.

history ($AVG_{noWC} = 3\text{min } 53\text{sec}$, $SD_{noWC} = 1\text{min } 07\text{sec}$, $AVG_{WC} = 1\text{min } 38\text{sec}$, $SD_{WC} = 18\text{sec}$). An independent samples t-test confirmed that the difference between the completion times was statistically significant ($p = .002$). Further, the subjects that used the WebComets tool have identified the correct time frame in 98% of the cases, while the group using the list-based browser history has reported a lower success rate with only a 71% average overlap of the detected collaboration time frame.

Additionally, the participants were given a post-task questionnaire inquiring about the usability, interaction and graphical aspects of the WebComets visualization. Moreover, participants had the opportunity to express their opinions freely about the application through a set of open questions. They perceived the comet tails as more enticing in terms of aspect. However, when working also with collapsed branches, their preferences shifted towards the beams as these would allow for an easier comparison of the various time intervals. The outcome suggested a high level of overall satisfaction with the visual representations (e.g., “I found the comets had useful information and were easy to understand”) and interaction possibilities supplied by WebComets (e.g., “I like the back and forward arrows. It’s a better solution than repeating the web page instance” or “I can oversee all important information in one view”), and that most participants would use such a tool for analyzing parallel browsing behavior and for detecting and comparing browsing patterns. Based on these findings and the results of the post-task questionnaire given to the participants, WebComets seems to address the issues raised by our requirement analysis.

3.4 Summary

In this chapter we presented WebComets, an interactive visualization for tabbed browser histories, a reflection of user online parallel browsing behavior. After highlighting the functional requirements, we focused on the interactive features and filtering capabilities of our tool. WebComets allows its users to more efficiently search for patterns in parallel browsing sessions by means of motif-based filtering as well as compare and analyze the tabbed browsing behavior of online users. Finally, an evaluation confirmed that our approach has met the initial requirements, and our users were able to quickly filter navigational information and detect patterns in the online sessions they were exploring.

Due to the value and complexity of the data generated by multi-user multi-session parallel browsing behavior and its suitability for collaborative exploration and analysis, the WebComets visualization was considered ideal for the testing and evaluation of our methods for supporting user-centered collaborative visualization. As a result, after a presentation of our techniques for supporting interpersonal awareness (Chapter 4) and collaborative interaction (Section 5.1), in Section 5.2 we introduce a touch-oriented adaptation of the WebComets visualization that allows us to evaluate the efficiency and effectiveness of our techniques in a collaborative information visualization application.

Chapter 4

Emotions in Collaborative Settings: Detection and Awareness

Contrary to what many people like to believe, emotions have been proven to affect almost every aspect of human existence. On an individual level, emotional states have been shown to affect creativity [89] and motivation [32, 60], as well as influence a person's judgement, decisions and problem solving ability [64, 102]. While all these processes can be closely connected with visualization, there are also studies about the direct influence of emotions on visualization scenarios [102, 165].

As such, the entire reasoning and analysis process that is specific to visualization experts can and is influenced by their emotional states and other subjective traits [152, 166]. Yet, the problem gains an additional level of complexity if we consider the highly dynamic social interactions that take place inside a group. Heer et al. [103] have already suggested the high importance of awareness and feedback in collaborative information visualization. However, as highlighted by Garcia et al. [93], awareness in terms of data and user activities might be insufficient for supporting an efficient collaboration and communication process. While these informational levels are certainly vital, they focus mostly on the awareness of system-user actions. However, the other side of personal interactions and their subtle dimensions should also not be neglected (e.g., user personality, intentions, emotional states, etc.). Moreover, both in terms of individual and collaborative visualization, measurements like user emotions and cognitive workload are ideal candidates for enhancing common qualitative evaluation techniques [5, 49, 210], which thus far have not always managed to capture the personal experience of the user.

After having presented our visualization approach for user parallel browsing logs in Chapter 3, this chapter revolves around two main topics meant to positively influence interpersonal interaction and collaboration: subjectivity recognition through a portable electroencephalographic (EEG) headset and the visualization of subjectivity in various environments and scenarios. More specifically, the structure of this chapter is presented in Figure 4.1, highlighting both the order of the tackled topics and their logical interconnection. The two inspected subjective dimensions are represented by user emotions and user insights, two related concepts that are deeply connected to social interaction and visualization effectiveness. As suggested in Figure 4.1, the visualization of user insights is beyond the scope of this thesis and will be address in future research.

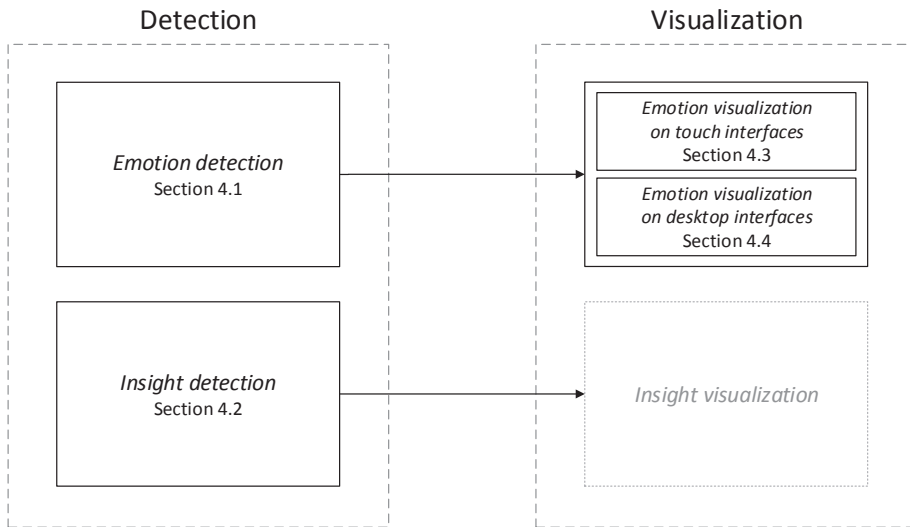


Figure 4.1: Structure and logical flow of the sections in Chapter 4. Note that the topic of emotion visualization has been subdivided in order to address the individual requirements of the two main interface types: multi-touch and desktop.

In Section 4.1 we explore the potential of a commercial wireless EEG headset to be used in the context of recognizing user emotional states in real-time and supporting the evaluation process. After validating the software framework of the device against other methods for reporting user emotions, we show in two evaluation scenarios how lightweight EEG headsets can be used successfully for non-intrusive affective classification.

In Section 4.2, we employ the previously gained knowledge to move the topic closer to information visualization, as we explore a potential correlation between moments of insight—a core concept in visualization—and user emotional reactions. Our results suggest that there is a strong connection between insight and emotions like frustration and excitement. We further argue that this indirect detection of insights opens the door for the objective evaluation and comparison of various visualizations techniques.

In Sections 4.3 and 4.4 we present two visualization frameworks aimed at increasing emotional awareness and aid system evaluation through emotion visualization. The first representation, called EmotionPrints, is capable of representing user emotional valence and arousal in the context of multi-touch systems, while the second one, called EmotionScents, offers a visualization of valence-arousal pairs specially designed for desktop graphical user interfaces (GUI) and standardized GUI widgets. With both visualizations, our goal is to enrich the interface with a new dimension of subjective information by adding support for emotion awareness as well as post-task analysis and decision making. In terms of real-time representations, this is achieved by showing user affective states in the moment when and at the location where the interaction occurs in order to increase affective self-awareness, support awareness in collaborative and competitive scenarios, and offer a framework for aiding the evaluation of touch and desktop applications through emotion visualization. Moreover, both visualizations have features that support offline analysis of user emotions, further enabling post-task analysis in order to explore user emotional reactions correlated with application features or social interactions.

The current chapter is based on a number of scientific publications. Section 4.1 focusses on the automatic detection of user emotional states with a commercial EEG device and is also described in “EEG-Based Measurement of Subjective Parameters in Evaluations” [47] and “Measuring Subjectivity – Supporting Evaluations with the Emotiv EPOC Neuroheadset” [49]. Further related to user emotions, Section 4.2 inspects the potential of detecting moments of insight through affective measures, a topic that was also covered in the publication entitled “Detecting Insight and Emotion in Visualization Applications with a Commercial EEG Headset” [44]. The last two sections focus on the visualization of user emotions on multi-touch and desktop systems, as presented in “Emotion Scents – A Method of Representing User Emotions

on GUI Widgets” [52] and “Emotion-prints: Interaction-driven Emotion Visualization on Multi-touch Interfaces” [53]. Details related to this last topic can also be found in the bachelor’s thesis “EmotionWidgets: Emotionen in GUI Controls” [273] by Kristin Suchner.

4.1 Detecting Emotions with a Commercial EEG Headset

The measurement of user satisfaction is and has been a core factor to introducing new goods and services to the market, in order to minimize the risk of a financial flop. Evaluation cycles are an essential step for developments, not only in the field of Human-Computer Interaction (HCI) or Computer Science. In fact, carefully designed evaluation tests are the decisive factor to find out if a solution (e.g., a new software product) or even a process (e.g., collaboration) will work as intended, based upon usability guidelines as well as user experiences. Still, often the outcome of an evaluation is somehow doubtful due to the inherent *subjectivity*. One major component of this user subjectivity is the emotional state which can not only influence evaluations, but also affect entire social and cognitive processes, including collaboration and decision making [64, 93, 102]. This makes it even more important to find real-time techniques for detecting and utilizing user emotional states in order to capture and raise awareness of this subtle dimension of social interaction and user experience.

With the approach presented in this section, we intend to capture user subjectivity in real-time, in the form of emotional states. For this, we utilize a consumer market electroencephalographic (EEG) device, the wireless Emotiv EPOC headset (Figure 4.2). We compare and validate the affective readings obtained through the BCI to widely-accepted evaluation techniques, like questionnaires, video-log analysis, etc. Our goal is therefore to employ a lightweight wireless EEG headset as a support device for subjectivity measurements, and even as a standalone system for fast first impressions.



Figure 4.2: The Emotiv EPOC EEG headset (images courtesy of Emotiv).

In the following, the focus goes towards validating the output of this EEG headset. Next, two evaluation scenarios are highlighted, in which the EPOC device is employed in order to capture in real-time the emotional states of each test subject related to his/her interaction with a given system. A discussion offers methods for optimizing the output of the EEG headset and reducing the error rates, while the last subsection highlights the impression of the users related to the potential of the device.

4.1.1 EPOC Headset Validation

The Emotiv EPOC device is a high-resolution wireless EEG headset that enables the acquisition and analysis of electric activity detectable on the human scalp. In combination with the corresponding EPOC framework, the headset captures electric signals that are mostly produced by the brain and classifies these in terms of expressions, emotions and commands.

Focusing on the technical features, the wireless EPOC headset is equipped with 14 electrodes (+ 2 references) that capture the electric signal generated by brain activity, as well as electrooculography (EOG) and electromyography (EMG). The 14 sensors require a saline solution for transmitting the signal and are spatially distributed according to the international 10-20 system for EEG measurements (AF3, AF4, F3, F4,




Figure 4.3: The 10-20 electrode distribution system and the locations of the EPOC’s electrodes [254]. The letters correspond to brain lobes—frontal (F), temporal (T), parietal (P) and occipital (O)—while the numbers further encode the positions—even numbers on the left, odd numbers on the right.

FC5, FC6, F7, F8, T7, T8, P7, P8, O1, O2, see Figure 4.3). The signal sampling rate is 128 Hz. An incorporated gyroscope captures the rotation of the user’s head, useful for the detection of gestures (e.g., nodding, head shaking). Additionally, the battery lasts for up to 12 hours, which is more than sufficient for most interaction scenarios.

Mainly a gaming device, the EPOC was designed with an entire spectrum of application fields in mind, ranging from interaction support for disabled patients to artistic expression and even market research. But, with features like affordability, non-intrusiveness and competitive precision [12, 272], this device has the potential of being successfully used in real-time evaluations of user’s subjective perceptions of

a multitude of complex software and hardware systems. Furthermore, compared to other widely used evaluation techniques, a wireless EEG headset has the advantage of offering results that are rarely influenced by external factors, as the main input data are the raw signals of the user’s brain activity.

In order to extract higher-level information from the headset without accessing the raw signals, the Emotiv framework and SDK supports the detection of a set of facial expressions (e.g., blink, raise brows, smile, etc.), emotional states (e.g., excitement, frustration, engagement, etc.) and cognitive commands (e.g., left, right, push, pull, etc., as well as user defined commands). While a step towards designing simple evaluation scenarios for capturing the users’ emotional states in real-time, the SDK is not open-source and the process of interpreting the various user states is not accessible to the EEG researcher. Still the potential of extracting emotional information from EEG readings has been shown in multiple instances (e.g., computing a user engagement index [20,218], real-time classification of facial expression via EEG [106], etc.). In fact, projects like Emokit¹ have tackled this issue partially by trying to reverse engineer the framework and offer users access to the raw EEG data on multiple platforms. However, there are still only limited studies about the accuracy of the EPOC framework’s output.

We argue that such a framework can bring multiple benefits to the field of computer science and subjective measurements, if it is validated. To overcome this problem and employ the EPOC device together with its framework for evaluations, we initially compare the data generated by the headset with widely accepted evaluation techniques (e.g., video log analysis, questionnaires) in a validation stage. While the EPOC device has functionality that allows the detection of facial expressions, emotional states and cognitive commands, we are only interested in the first two areas in our tests. This is due to the fact that expressions and feelings are mostly subconsciously activated and of high relevance for the overall subjective perception of the users, while cognitive aspects are more related to the conscious control of a PC-related system via a reproducible brain signal configuration.

¹Open source driver for accessing raw data from the Emotiv EPOC EEG headset, <http://github.com/qdot/emokit> (February 2014).

In the following validation and evaluation scenarios, the EPOC’s capacity to accurately detect and classify facial expressions and emotional states, as well as its employment in the detection and classification of user subjectivity. Since the completion of our validation, described later in this section, other research has also focused on evaluating the EPOC’s capabilities in terms of facial expressions [106], emotional states [129, 194], cognitive commands [277, 312] and other specific EEG attributes [12, 75], with overall promising results.

Our study has been executed on 12 test subjects. These participants had a diverse cultural background and a varying level of familiarity with computers, thus representing a broad spectrum of potential users. Most of them had little to no previous experience with EEG systems, while some were previously subjected to EEG medical tests. Nonetheless, the concept of an EEG headset in Computer Science was not familiar to any of the subjects. The test group was composed of four women and eight men, aging 21 to 52, with an average age of 29.75 years.

Facial Expressions

Initially, we focused our attention towards the facial expressions suite incorporated in the EPOC framework, as facial expressions have the capacity of reflecting the emotional states of a person. Nowadays, there are multiple techniques for capturing expressions, and the majority of these rely on video recordings of the user’s facial gestures together with the employment of a classification methods. Still, optical methods restrict the user’s mobility, as the camera(s) must record the person’s face at every moment. Also, many for real-time facial expression recognition algorithms still have difficulties with minor contractions of the facial muscles (i.e., microexpressions) that are usually detectable for a human observer, but often missed by a classification algorithm. Therefore, based on previous research [13, 187] that suggests that EEG/EMG signals incorporate head muscle movements, we considered capturing the facial expressions of users with the EPOC headset. As the Emotiv framework includes methods for detecting facial gestures, we compared the output of the headset with the results of a video log analysis in order to detect the accuracy of the EPOC detection and interpretation.

In a first step, the EPOC device was mounted on the test subjects, ensuring the correct distribution of the 14 electrodes on every user's scalp. Afterwards, the participants were granted a ten-minute accommodation phase in which they would familiarize themselves with the feel of the device. Furthermore, this period of time was used for an initial calibration of the detection weights that are part of the framework. Details about this process will be highlighted in Section 4.1.3, focusing on good practices and error reduction when employing the EPOC device.

Each user was asked to watch a computer screen on which words would appear that represent facial expressions (e.g., smile, blink, look right, etc.). For the couple of seconds that the words would be displayed, the users were instructed to execute the corresponding facial expression for as long or as often as they feel comfortable, but only while the words are still visible. Our tests involved 8 facial expressions (Figure 4.4), where each expression would randomly appear on the screen at least three times.

In parallel, a webcam was recording the face of each test subject. After the session was completed, the video logs were analyzed and transcribed in detail, thus enabling a comparison with the information given by the headset. The transcripts that were derived include the interpretation of facial expressions as well as any verbal communication between the participant and the researcher. Furthermore, the transcript is complemented by comments about expressions and body language as well as pauses in activity, if considered important. The ultimate goal of the evaluation was to obtain both quantitative and qualitative information about the validation process.

The results obtained for the facial expressions are presented in Figure 4.4 and show that the correct detection varies between 70-100%, depending on the particular facial expression and the detection rule. The simplest rule (indicated by the blue bars in Figure 4.4), which checks if an expression was recognized by the device during the available timeframe, seems to give the best results with a detection rate of 82-100%. The algorithm with the highest error rates was the one that computed the length of all detected expressions during the timeframe and checked which represented the longest one. Still, even in this case, the recognition rates are between 70-91%. These results are put into perspective by the fact that some participants executed certain facial expressions very differently from others, as suggested by the comments added to the video log.




Figure 4.4: Percentage of correctly recognized facial expressions (from left to right: look left, look right, raise eyebrows, blink, wink left, wink right, smile, and clench) with the EPOC headset. The detection was executed by three algorithms (left to right): was the executed expression reported at all (blue); was it the longest reported expression during the corresponding period (red); and was it the most frequently reported expression during the corresponding period (green).

Emotional States

In a second stage of our investigation, we inspected the accuracy of the emotion detection suite supported by the EPOC device. As emotions are at the core of the subjective perception in humans, establishing the capability of this headset to directly detect emotions would be vital for our research. Again, as the framework that offers these features is proprietary, we have to verify the correctness of the data output for the headset. To achieve this, the test subjects were asked to participate in a series of short tasks that were selected specifically to induce a certain emotional reaction. After each task, the users would have to complete a short questionnaire that would represent the term of comparison for the output of the EEG device. A set of similar results between the data generated by the device and the answers of the users to the questionnaire would allow us to consider the EPOC headset as a viable alternative to common techniques of evaluation.



Figure 4.5: Sample video frame from the user calmness measurements.

The tasks the users had to complete had a length of maximum three minutes. Similarly to [129,194], these tasks involved watching a video clip, listening to music, and playing dexterity and memory games on the computer. Each task was selected to generate a certain emotional reaction, and only this emotion was inspected by the headset. Note that the users were not informed about the inspected emotion previously to the task. In order to additionally confirm that the task materials (video, music, games) were appropriate for generating the expected emotion, each post-task questionnaire inquired about the general potential of that material to awaken the corresponding feeling; the results in this case were 100%, as all subjects found the materials adequate for inducing the selected emotional response.

The selected emotional responses were calmness, meditation, engagement in dexterity and mental tasks, and excitement. The tests were ordered from calmest to most exciting, as arousal states often have longer recovery times [68,107]. “For example, following a fear-provoking encounter, some individuals show a persisting heart rate elevation that might last for minutes, whereas other individuals show a comparable peak and rise time, but recover much more quickly” [68].

Between each task, there was a short break. During the tasks, the device returns a constant stream of values at an average frequency of 4-5 Hz for each of the emotions mentioned above, instead of a simple Boolean expression. These normalized values encode a gradient along which the emotional activation is distributed. This meant that, as questionnaires usually employ scales, we had to find a classification of the emotional intensity for each detected affective dimension. Additionally, the data has a different baseline for each emotion and each user. In our tests, each participant's baseline was composed of relatively stable values offered by the framework for each of the detected emotions. This affective baseline was established during the task intermissions, when no stimuli were offered to the participants, they did not engage in any activities and they reported no particular state of arousal. As a result, our measurements focused on interpreting gradual and sudden fluctuations in the reported affective states, instead of quantifying different response levels.

More precisely, each emotion had a classifier attached that would distribute the intensity on a 5-point scale. This was achieved by implementing an algorithm that would in real time inspect the value vector, and compute the tangent and the local distance between the current and the minimum value. Usually, this minimum value would represent the baseline established during the calmness tasks. The algorithm implemented a set of four triggering conditions able to classify the intensity of each computed emotion. Two of the triggers were computed based on the current tangent to the data stream; these triggers were considered activated at angles larger than 30, respectively 60 degrees. The other two triggers focused on slow, but substantial deviations from the established baseline or minimum. As such, an emotion would be considered as not triggered at all if the current values are close to the baseline and none of the four conditions are satisfied.

Similarly, the possible answers included in the questionnaires were distributed on a 5-point Likert scale [35, 161], in order to allow a simple comparison with the output of the EPOC device for each emotion. The questionnaire for each task would inquire if the participants had experienced any of the considered emotions. Consequently, they had the possibility to express their option by selecting one of the following: strongly disagree, disagree, neutral, agree and strongly agree (interpreted as 0, 25, 50, 75 and 100%).

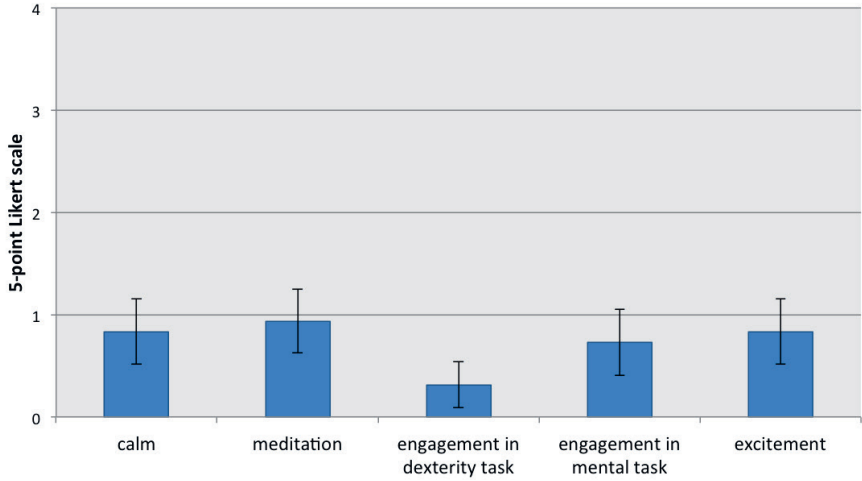


Figure 4.6: Average difference between the EPOC device output and the questionnaire results for the emotional evaluation (left to right: calmness, meditation, engagement on dexterity task, engagement on metal task, excitement), with both on the same 5-point system: 0, 0.25, 0.5, 0.75, and 1.

The results of this comparison are highlighted in Figure 4.6 that represents the average difference in percentage between the questionnaire answers and the EPOC data. More precisely, an average difference of 1 means that, on average, the difference between the user’s answer in the questionnaire and the EPOC output was one Likert interval out of four. For instance, if the user’s answer was “strongly agree” and the device detected “agree”, the difference is one interval. Moreover, the higher difference values for the meditation task might suggest that the understanding of concept of meditation is highly subjective.

As emotions do not always have an external manifestation, we could not verify the output of the EPOC device directly like we did in the case of the facial expressions, where we compared the data generated by the headset to the video logs. Therefore, besides computing the standard deviation for the average differences (Figure 4.6), we validated the results by executing a paired samples t-test for the EPOC and questionnaire data gathered during each task. For these paired sets, no significant difference could be detected, thus suggesting that the EPOC headset has the potential of computing accurate interpretations of human emotional states.

4.1.2 Emotional Response Classification with the EPOC Headset

Based on our positive findings for facial expression and emotional state detection and classification with the EPOC device, in a next step of our research we focused on real-time classification of affective states for the users of complex software systems. To achieve this, we proposed two evaluation scenarios that would have the capacity to generate multiple emotional responses. More precisely, the emotions that were considered in our tests were engagement, excitement, satisfaction and frustration. As the EPOC framework does not directly compute satisfaction and frustration, we considered a combination of the levels of excitement and engagement with various facial expressions (e.g., smiling, clenching teeth). Similarly to the validation of the emotional suite described in Section 4.1.1, we compared the values the headset returned for these emotions with the results of a questionnaire completed by each participant.

In the first scenario, we evaluated the emotional effects a spot-the-difference task would have on the test subjects. For this mostly mental/visual task, we sequentially presented the users with three pairs of slightly different images. In a second scenario, the users had to play a first-person shooter game (FPS) with a competitive setting, where shooting the adversaries was not the only way to win. For this dexterity task, we also offered an initial introduction and suggestions, as some users were not familiar with the type of game and interaction.

As both scenarios had the potential of generating an entire set of emotional responses distributed over a longer period of time, this implied that in the post-event questionnaire users might find it difficult to remember their inner states at a particular moment. To compensate, we considered slightly different approaches for the two tasks. In the spot-the-different task, the final questionnaire only inquired the users about the presence of a certain feeling during the scenario (e.g., “Did you at any point during the task feel frustrated?”). In our view, this reduced the memory load of the users, as now they could disregard the temporal aspect and the exact coupling of an emotion with an event.

Contrary to this approach, in the first-person shooter game a task supervisor had the ability to pause the game whenever he noticed a significant emotional spike returned by the EPOC headset. During these short breaks, the supervisor would mark the time of the pause and pose the questions from the questionnaires (about all four emotions) to the users. To avoid bias, the task supervisor would also randomly pause the game at points where no particular emotional response was generated by the EPOC device. After the game was restarted, the users were granted a re-adaptation period in order to allow them to focus on the game again. During this two minutes period, the FPS game could not be paused again by the task supervisor.




Figure 4.7: Average difference between the EPOC device output and the questionnaire results for the two scenarios: the spot-the-difference task (left) and the FPS game (right). The considered emotional states for each task are (left to right): engagement, excitement, satisfaction, frustration. The same 5-point Likert scale from Figure 4.6 applies.

Figure 4.7 presents the results of our two evaluation scenarios in terms of average differences between the EPOC outputs and the questionnaire answers plus video log transcript. Note that the highest average difference (1.32 Likert intervals) was detected for the excitement level in the first task, while the lowest (0.32 Likert intervals) for the frustration level in the second task. Also, most other resulting average differences are in the scope of 0.7–1.0, similarly to the results presented in the testing phase in Section 4.1.1. The satisfaction and frustration detection yielded the most similar

results to the ones of the questionnaire in case of the FPS game, with an average difference of around 0.4. This might suggest that, even if disruptive, asking certain questions during the tasks gives users the capacity to more accurately estimate and remember their emotional states, compared to the effort of recalling them after an entire set of emotional experiences.

To further validate our results, we again computed the standard deviation for the average differences (also visible in Figure 4.7), and obtained a maximum of $SD = 0.48$. We also executed a paired samples t-test for the headset data and the questionnaire results for each emotion. All cases showed no significant difference, except for the excitement detection for the first task. One could speculate that this is related to the users' misconception about excitement in mental tasks, a fact further suggested by our open-end questions at the end of the entire evaluation.

4.1.3 Limitations and Error Reduction in Emotion Classification

All complex EEG systems used to detect brain activity are sensitive to a set of elements that can negatively influence the accuracy of their measurements. These influences, in terms of interpretation of higher-level facial, emotional and cognitive aspects, are manifested as false positives and false negatives. The detection of states that did not take place in reality is called false positives, while false negatives refer to the system's incapacity to capture an existing state or command. These undesired results are not only strongly connected to the capacities and accuracy of the BCI system, but also to the training level and state of mind of the users, as well as to the physiological and extraphysiological aspects of EEG. This triad—BCI system, user and surroundings—has to be optimized in order to reduce any influencing elements that disturb the measurements and their interpretation. In the following, we want to highlight a set of basic suggestions for improving the performance of the EPOC BCI headset based on previous EEG research, technical aspects of the EPOC device and our experience with this system.

Of central importance to any brain activity measurement is the user of the EEG system. In order to supervise and improve the results of BCIs, one has to better understand the neurological and physiological aspects that directly influence the electrical signals generated by the human brain and how these are captured by any EEG system. Most EEG devices actually capture a mixture of skin, muscle and nerve activity instead of a pure signal generated by the electrical activity of the neurons [290]. These contributions to the EEG signal are usually not accessible independently [132], thus resulting in a system that actually reflects the summation of the synchronous activity of different biological systems. Moreover, one has to note that EEG measurements only capture the electrical signal generated by the neurons that have a radial orientation to the scalp, and not the ones that are tangential.

In terms of physiology, one has to consider that artifacts generated by muscles movement should in general be avoided, unless the measurements are aimed at classifying facial expressions. In this case, eye movements or signals from the frontalis and temporalis muscles are especially useful, as they are rather easily distinguishable in the processing of EEG/EMG data [19].

While many EEG systems are heavily shielded against unwanted external influence, the test environment still represents an important factor in determining the levels of noise that is injected into the signals. Such extraphysiologic artifacts [19] can be generated by interference with electrostatically charged objects (usually other people in contact or close to the user), high-frequency radiation (electric communication devices - radio, cell phones, etc.) or EM field generators (powerful electrical appliances). All these influences should be considered when designing a scenario and the positioning of the test subjects.

Nonetheless, EEG measurements would not be possible without the corresponding devices, capable of measuring these small electrical signals on the human scalp. For this component of the triad too, there are a couple of suggestions that may improve the overall results of the measurements. One of the problems that many mobile EEG devices still face and which is particularly important for collaboration and visualization is the level of comfort they need to offer during their usage. Being involved in a collaborative as well as a visualization task might require the user to wear the device for extended periods of time, where the level of comfort can make an impact on the process. Furthermore, while wireless systems, such EEG headsets are still susceptible

to motion, especially to fast head and body motions. In this sense, the electrodes and their positioning represent a vital element for accurate EEG measurements. While the EPOC headset almost implicitly positions the electrodes on the human scalp based on the 10-20 system (Figure 4.3), the users/researchers still have to pay particular attention to reading and following the positioning instructions included in the product manual. Furthermore, the electrodes should have sufficient saline solution applied to them and they should have a good contact with the scalp. The level of contact and the power of the received signal are highlighted by a feedback system included in the Emotiv framework. When working with a non-dry sensor setup, like the one the EPOC uses, one has to consider that the solution that supports the transmission of the electrical signal from the scalp to the electrodes can slowly evaporate. In our tests, the wet saline sensors could only be used for short periods of time (compared to gel-based methods) of up to two hours before the signal strength started to degrade and the sensor pads needed to be remoistened.

Specifics of the EPOC Headset

Besides physical aspects, there are elements specific to the EPOC's framework that can be considered for improving the rate of correctly detected emotional states and facial expressions. These methods revolve around a key element, highlighted in multiple publications employing the EPOC system: training. Still, these research projects mostly focus on the conscious cognitive control of various devices via the EPOC interface. As at the core of the framework are various Machine Learning algorithms, which suggests that users have the possibility of improving the sensitivity of the detections by training new samples for the commands they wish to execute.

On the other hand, there is a set of specific aspects of the training process that apply for the facial detection and emotion suites, where consciously offering training samples is rather difficult. For facial expressions, users have the possibility to fine-tune weighting controls for the detection of each expression, thus employing a trial-and-error method. The sensitivity of the detection is vital as each subject has a slightly different positioning of the headset and different ways to control their muscles for generating a gesture.

To highlight this in our tests, we executed a second round of validations for the facial expressions after another 20 minutes of training, in which the supervisor together with the user would tune the sensitivity of the system. Similar to the first stage, the test subjects were not told to change their natural way of expressing different facial gestures, not even in cases where the facial expression executed in the first stage was hard to classify for human observers. Again, the subjects executed all facial expressions in random order, each expression at least three times. As a result, we noticed an increase of 12% of the correct detections for facial expressions. Nonetheless, based on our tests, there seems to be no reduction of these errors through learning or practice.

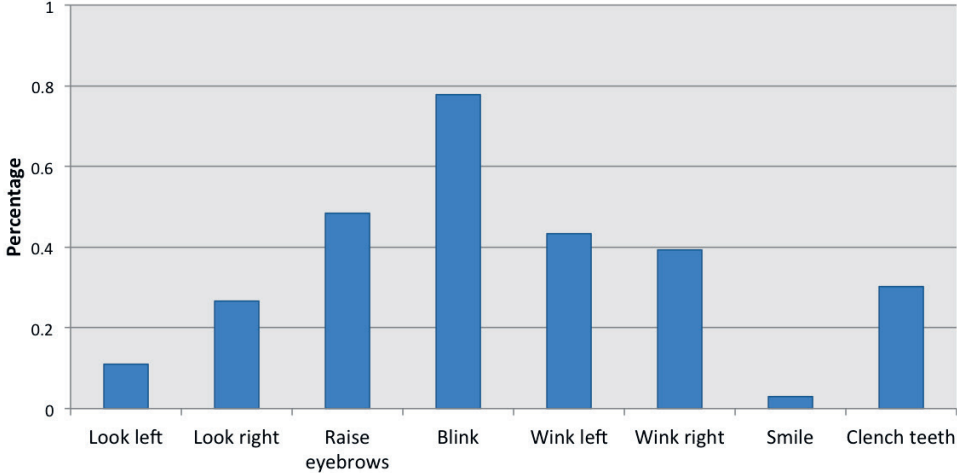


Figure 4.8: Percentage of reduction for the false positives when considering groups of facial expressions.

After further exploring the generated data and considering the theoretical aspects, we noticed some patterns of similarity in the false positives. As false detections were mostly incorrect detections of other expressions, we considered grouping facial expressions by the group of muscles that are activated in order to execute them. We categorized looking left and right into eyeball movements; raising eyebrows, blinking, winking left and right into frontal muscle movements; and smiling and clenching into jaw muscle movements. Finally, we explored the level of false positives when considering these groups of expressions and not every individual expression—e.g., if the blink would be detected as a wink left it would be still accepted as a true positive,

as both expressions involve frontal muscles. The results are summarized in Figure 4.8. One can see that for each facial expression the reduction of the false positive rates is significant, peaking at 78% for blinking. This suggests that the EPOC device is proficient in the detection true and false positives for groups of muscles, but the classification of the exact muscles motion can sometimes be inaccurate. Knowing this, evaluators can improve on the detection of feelings in users by not only relying on facial expression, and further accessing the capabilities of the emotional suite.

For the case of detecting emotional states, it is difficult to define true and false positives in terms of the EEG signal or the EPOC data stream. This is due to the fact that, similarly to our test, affective states detected with such a BCI device can only be compared with perceived emotions that the user is also aware of. Still, it seems that establishing a baseline for each emotion and each subject, as well as the usual deviation levels from this baseline, can improve the detection rate. Furthermore, in our experiments, many false positives were generated by short peaks (under five seconds) in the data stream corresponding to a certain emotion. This is most likely influenced by the fact that the EPOC framework uses complex computations for the emotion suite (also suggested by the lower output rate, around 4 Hz), computations that involve inspecting an entire set of previously acquired raw EEG values.

User Feedback

After the completion of all the tests and questionnaires, the users were asked a set of open-end questions about the device, its use and their mental state during the entire process. The overall feedback of the users with respect to the EEG device was positive. Additionally, after a short accommodation period (a couple of minutes) most users felt comfortable and able to perform well with the headset on.

Still, in terms of practicality and applicability, there are some limitations that need to be considered. Negative feedback was given for the limited wearing comfort of the EPOC headset over a longer period of time, as it becomes inconvenient after about 30 minutes. In some cases, this can even result in pain, as expressed by some testers (e.g., “My head hurts”, “Is my head supposed to hurt from this?”). This is due to the side-pressure that the device applies on the head in order not to move on the

scalp, as this would disrupt the EEG measurements. Improving the hardware itself in next generations is a possible solution to overcome this issue. Nevertheless, currently it seems necessary to carefully design the evaluations to allow regular breaks, as the device becomes uncomfortable to wear after some time.

Also, two users complained about the inconvenient non-dry sensors (“But it dirties my hair”). While the discomfort of the users is real, the sensors do not represent a hygiene issue themselves, as the liquid used for the sensor is actually a disinfectant. Other devices like NeuroSky², Mynd³ and EmSense⁴ represent an attractive alternative for EEG wearable systems as they use dry electrodes, but at the same time, these systems have to face additional challenges for capturing the small electrical signals [229].

One participant could not be evaluated, as this particular candidate had “springy” hair and the EPOC device sensors could not make a good contact with the subject’s scalp, resulting in very weak electric signals. Also, with three of the 12 subjects, we had pronounced difficulties in positioning the EPOC device, due to the length of their hair. Eventually, the position was good and the signal strong, but it took several minutes of adjustments. Ten subjects clearly stated that they did not perceive the device consciously once the tasks started and that, in their opinion, the presence of the headset did not influence the way they acted.

Finally, we considered it relevant to inquire the users about their state of mind on the day the tests took place, as an overall positive or negative state could influence their capability to concentrate, as well as affect their emotional output. Most subjects considered that they had a normal day, while two felt it was “a bit stressful”. Still, no significant variation was observed in terms of results between these two groups.

4.2 Detecting Insight through User Emotions

Insight, epiphany, eureka moment, Aha! effect [158]—these are all names for one of the most intriguing and even mysterious process through which humans gain knowledge. But what is insight really and how does it enrich our capacity to gain and manage knowledge? There are many definitions, each capturing a slightly different aspect of

²NeuroSky, <http://www.neurosky.com> (February 2014).

³Mynd, <http://www.neurofocus.com> (February 2014).

⁴EmSense, <http://www.emsense.com> (February 2014).

the experience. The Merriam-Webster dictionary defines insight as “the act or result of apprehending the inner nature of things or of seeing intuitively”. Encyclopedia Britannica exposes it as the “immediate and clear learning or understanding that takes place without overt trial-and-error testing”. Whatever the definition, all suggest the presence of a moment of extreme clarity, a moment when a solution is found that satisfies all conditions for the problem that is inspected.

While this concept has been around for hundreds of years, it has been only introduced in psychology at the beginning of the last century [34], as the German term “Aha-Erlebnis”. Since then, the process of insight has been investigated from the perspective of many fields, like medicine, cognitive neuroscience and computer science, to name just a few. At the same time, some researchers dislike any reference to spontaneous Aha! moments because it suggests irrationality. Still, many of world’s most famous discoveries have been achieved by people experiencing a moment of epiphany. Isaac Newton claimed having a moment of clarity when he observed an apple falling from a tree, an insight that led to the formulation of the theory of gravity. Similarly, Friedrich August Kekulé von Stradonitz experienced the ring-like structure of benzene in a daydream [168].

Besides the purely knowledge-related aspects of insight, particular experiences suggest that moments of epiphany are sometimes accompanied by extremely powerful emotions [148], like the joy of understanding a problem or the excitement of decoding a riddle after a timely process of analysis. These moments of triumph have in many instances shown their potential to shift the emotional states of a person. Still, “the shock of recognition” is not always a side effect of the Aha! experience [207], and further investigation is required to establish a possible correlation between insight and emotion on insight.

Furthermore, directly detecting moments of insight is difficult, and neuroscience has struggled to capture these events in real-time. While modern methods like fMRI scans support the identification of Aha! moments [55], these approaches are still very restrictive and even intrusive operations for the subjects. More importantly even, insights are one of the core elements of visualization. One of the researchers that heavily influenced this field has also supported this clearly by stating that “the purpose of visualization is insight, not pictures”. Nonetheless, adjacent processes like emotional reactions generated by the excitement and joy of insight might be more simply de-

tected by mobile brain-computer interfaces than moments of insight themselves. BCI devices can represent the key for a less intrusive, indirect identification and observation of periods of insight, as well as a migration of insight detection to wherever it takes place without limiting the environment of its existence, i.e., medical facility. Perhaps more importantly, this indirect detection of insights—be it in visualization or in general—can support the evaluation of a system or process by capturing its perceived user efficiency.

In the following, we shortly highlight related work in the field of insight research. Next, a preliminary study is presented that involves the observation of brain signals by the EPOC headset and the translation of these signals into emotional reactions generated by moments of insight. We highlight the results of this study, and capture some advantages and limitations of indirect EEG detection of insight-related patterns.

4.2.1 Insight and Visualization

Many scientific areas have taken it upon themselves to bring clarity to the concept of insight. As a result, various characteristics of insights have surfaced during the past years, some more relevant than others in comprehending the series of processes that converge to create an Aha! moment. Studies have found that insight can be seen as a two-phase process [222]. During an initial step, a subject tries to systematically explore the space of possible solutions to the task. If this approach fails to give results in a certain time-frame, an impasse is achieved that can manifest itself in the form of frustration [240]. People try to overcome this impasse in a subconscious manner that builds upon relaxing the constraints of the problem or approaching it in a non-conventional manner (thinking out of the box). If the change in the mental representation of the problem is successful, the second phase is reached, the impasse is overcome, and the subconscious process suddenly and unexpectedly provides the person with a piece of information—an insight.

Studies suggest that the presence of prior knowledge about the problem or tasks as well as knowledge of one or multiple possible solutions or patterns, can interfere with the unconscious processing that leads to an insight [9, 314]. The reduced prior knowledge only adds to the unpredictability of this concept, which is one of its essential characteristics derived from the complexity of mental activities. In [178], insights are

considered in terms of pattern matching, where the mind is trying to establish an approximate fit between the set of current patterns and previous experiences. Further, a categorization is highlighted involving the information content of the epiphany in terms of anticipation: to recognize (expected information) and to notice (unexpected information).

Besides the field of psychology, various studies from medicine and cognitive neuroscience have focused on pinning down the processes and brain activity in the moment of insight. Most of these employed brain-imaging technologies, like electroencephalography (EEG) [133, 151, 256] and functional magnetic resonance imaging (fMRI) [28], to observe the brain activation patterns of subjects while solving a wide range of insight-connected problems. Participants were asked to solve specific insight problems, visual and logical riddles [167], and anagrams [11]. Some of these problems, like anagrams, are used because their solution can be achieved in at least two ways: through a conscious, systematic search of the space of possible solution or through sudden insight that appears abruptly in the conscious mind [28, 150]. The experiment results suggested that tasks that involve problem solving via insight activate certain parts of the human brain [27, 133, 150], leading to the possibility of detecting when a subject experiences an Aha! moment, and distinguishing this from simply finding a solution based on a systematic search.

But what about fields like information visualization that have the concept of insight at their core? Over the years, researchers have focused on defining insight and its importance for visualization [55, 201, 216, 239]. Most famously, insight is defined as the “purpose of visualization” [39], the ultimate goal by which successful representations and interactions should be measured. But how can we measure something as unpredictable and multifaceted as insight?

Most approaches in the visualization community try to achieve this by including characterizations that are aimed at defining insight in an objective, quantitative manner [201, 239], with attributes like time, complexity, domain value, depth, uncertainty, unpredictability, correctness, expectedness, and others. Attention is sometimes focused to a particular topic, like cartography [168], to investigate the appearance of insight when working with a certain type of representation.

Still, in many publications, the focus quickly shifts towards the importance of insight for evaluating visualizations. If insight is the purpose of all visualization, then it should also be the measure by which the quality and functionality of visualizations is determined. Currently, this is achieved in most cases by performance and accuracy experiments on restrictive benchmark tasks. Sadly, such restrictive tasks often introduce bias or capture only the performance for a particular type of task without giving answers about the performance of another. While [201, 239] highlight viable alternatives to this by suggesting open-ended protocols together with a set of quantitative measures for insight, such experiments could represent an intrusion in the analysis flow of the user by introducing interruptions or imposing the think aloud method (i.e., concurrent verbal protocol or CVP).

In the following, we highlight an approach that overcomes some limitations of the previously presented methods by relying on the EPOC mobile EEG solution for detecting moments of insight during visual problem solving.

4.2.2 EEG Detection of Emotion and Insight

As moments of insight are accompanied by powerful emotions of joy and satisfaction on discovery or comprehension, the question arises if an objective connection can be established between the Aha! moment and the experienced emotional states. In order to evaluate if insight generates emotional reactions that are detectable by means of EEG measurements, we devised a preliminary experiment that focuses on capturing the emotions of users while involved in visual problem solving.

In this study, we build upon the validation of the EPOC device and its capacity to detect estimate states—as presented in Section 4.1—to explore the existence of a correlation between insight and emotion. More precisely, the spectrum of emotions that is considered in the following experiments involves only the excitement and frustration levels of the participants. The ultimate goal of this endeavor is the analysis of how well emotional states can reflect the presence of insight, and if capturing these states by EEG enables the detection of Aha! moments in information visualization techniques. The final purpose of this is not only to aid the evaluation of visualization approaches, but also to unveil a novel dimension of subjectivity in collaboration.

The current study involved six participants with a good knowledge of visual representations and visualization techniques. The subjects were given a set of four tasks, two represented by visual insight problems and two information visualizations. For the visual riddles, the subjects had to find a unique solution, most likely resulting in a single fundamental insight. This allowed for a simple comparison of the moment of insight with the emotional states prior and during the discovery. At the same time, for the visualizations the participants were asked to find as many insights about the data as possible. For each task, users had ten minutes to offer their insights.

Insights take time to process and clarify in the mind. Carl Friedrich Gauss said once after an epiphany: “I have the result, only I do not yet know how to get to it” [74]. Therefore, once a user would report an insight, the EPOC output before this moment was inspected. More precisely, fluctuations in the levels of frustration and excitement were inspected for a time period of up to two minutes before the insight. In this context, a binary classification was used that would simply fire a rule when the corresponding readings for frustration and excitement passed a predefined threshold.

Visual Insight Tasks

For the visual riddles, all participants were initially subjected to a larger set of problems, of which only two were selected—Eight Coin Problem (Figure 4.9) and Matchstick Arithmetic (Figure 4.10)—that none of the subjects reported to know beforehand. For these two problems, only in 58% of all cases a solution was reached. In other words, the six participants reached an insight in 7 cases out of 12. Figure 4.11 highlights the correlation between insight and emotions in these cases.

Our results allowed us to notice that over 80% of those who managed to solve the visual riddles have felt frustrated in the two minutes before the insights. This is also suggested by other publications, that cite frustration or deadlock as a prerequisite for the generation of an Aha! moment [168]. In a slightly lower percentage of cases, the subjects have also experienced excitement in the seconds prior to the insight.




Figure 4.9: Representation of the eight-coin problem: the initial configuration for the coins (left), the solution to the problem (right). The configuration of the coins has to be modified by moving only two coins, such that in the new grouping each coin touches exactly three others [205].




Figure 4.10: Example of a matchstick arithmetic problem. The participants would be given an initial configuration (top), and they would be asked to correct the arithmetic statement by repositioning only one matchstick (bottom).

While these results by themselves give us a reduced amount of information about the connection between insight and emotion, Figure 4.12 captures the percentage of emotional reactions for subjects that have not experienced insight at all. The lack of insight for these participants was suggested, on one hand, by their lack of a solution for the problems, but also by a post-task questionnaire that each of them filled out.




Figure 4.11: Measured emotions with the EPOC headset in the presence of insight. The bar chart presents the average percentage of cases where (left to right) frustration was detected before insight, excitement was detected during insight, and both frustration before and excitement during insight were measured.

By inspecting both Figures 4.11 and 4.12, one notices that in both cases the frustration levels are around 80%, independent of the presence of an insight. But at the same time, the detection of excitement is much more likely in subjects that have experienced an insight. When looking at both of these emotional states, excitement and frustration were detected for 72% of the experienced insights. At the same time, the combination of excitement and frustration only appears in about 20% of the cases where no insight was gained by the subjects.




Figure 4.12: Measured emotions with the EPOC headset in the absence of insight. The bar chart presents the average percentage of cases where (left to right) frustration was detected and not followed by insight, excitement without insight, and the presence of both emotions when no insight was achieved.

As frustration seems to be a recurring element in problem solving, the time of appearance for the feeling of frustration was analyzed. Our results suggest that states of frustration tend to be measured in participants more often during the later stages of an exploration process (second part of the ten minutes window). Also, emotional states of frustration that last longer (over one minute) are more likely to be followed by excitement, which we hypothesize might be correlated with insight. Further research involving more tasks and participants will be required to confirm these results.

The two visual problems were followed by a questionnaires related to the emotional and mental states of the participants. After explaining what an Aha! moment is, we asked those that reported answers to the problems if they experienced an insight in this sense, or if it was a systematic search-and-find process that generated their answers. All the participants that experienced frustration and excitement, and that also reported the solution to the task, have confirmed that they experienced a moment of insight. On the other hand, in two instances, participants that supplied a solution and reported experiencing an epiphany were not reported by the EEG device as experiencing an increased frustration and excitement level.

Visualization Tasks

For the information visualization tasks, we selected two visualizations from the ManyEyes website, as it harbors various datasets represented by widely accepted visualization techniques [293]. More so, as the visualizations are collaboratively explored and annotated, one can detect those that have a popular thematic and a high potential for revealing patterns and supporting hypotheses manipulation. The popularity of the visualizations was important in the selection process, as it could suggest the overall effort that users would invest in mining for insight in that representation. At the same time, a visualization that captures the tendencies of a topic that is highly relevant to the analyst has, in our view, a higher chance of generating an emotional reaction.

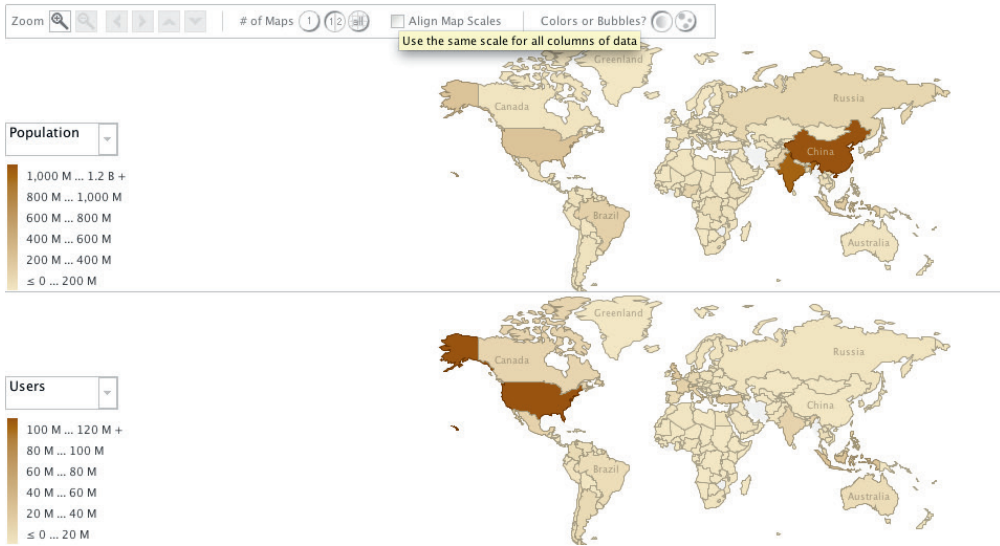


Figure 4.13: ManyEyes map visualization employed in our experiments.

The two visualizations that were selected contained data about global demographics and social media, and were represented as a stacked graph and a cartographic visualization, respectively (Figure 4.13). The participants had an accommodation period with the ManyEyes website, during which the main supported interactions were highlighted to them. Before beginning the task, the test subjects were instructed to find all possible insights in the visualization. This approach is similar to the one in [201], where insight was also observed by applying an open-ended protocol, without additional restrictions to the insights that were considered.

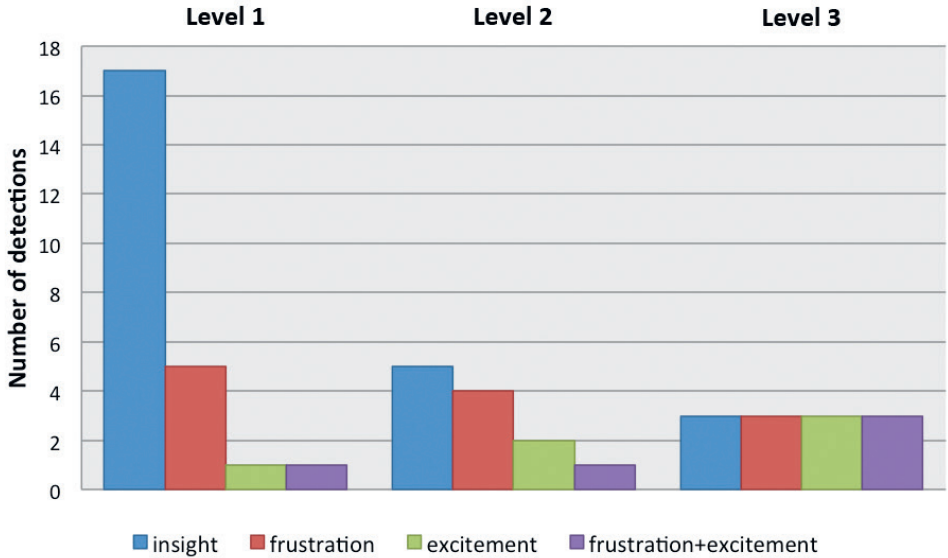


Figure 4.14: Correlation between the number of insights and the instances where frustration, excitement and frustration-excitement pairs were detected. The results are grouped by depth of insight: the four leftmost bars represent the values for depth Level 1, the next four for depth Level 2, and the last four bar encode the number of insights and corresponding detected emotions for the deepest insights.

Furthermore, it was also suggested to the subjects to focus more towards deep insights that involve new hypotheses and multiple data types, to avoid noticing only simple facts about the data. Similarly to [201] and [239], all spawned insights were grouped by depth into three levels: the first level refers to trivial insights that include direct observations of one data type; Level 2 insights that are generated by a combination of multiple data types or insights about a process; and Level 3 insights that refer to new hypotheses about the underlying information. The EPOC headset was used to inspect the levels of emotional frustration and excitement during the interaction of the users with the visualizations.

Figure 4.14 presents the correlation between the number of generated insights and the various emotional states the users experienced. The bar chart is divided into three sections, representing the different depths of the captured insights and their corresponding emotions. The number of simple insights seems to dominate the representation, as deeper insights were more rarely detected. This fact is even more visible in Figure 4.15, where every single eureka moment and emotional state was plotted along a time axis.

Although the number of deeper insights is lower than the one of trivial observations, one notices the fact that deeper insights have a higher probability of generating an emotional response, especially a higher probability for excitement during the Aha! moment. This culminates in our tests for the deepest insights with a detection accuracy of 100%, via the EEG measurements of emotions, when considering both the prior experience of frustration and the excitement on discovery. Note that in Figures 4.14 and 4.15 the results of the two visualizations are convoluted, as no significant differences could be detected between the results for the two visualizations.

By using the temporal dimension, we also notice in Figure 4.15 that users more quickly detect the simpler insights than deep ones, and that the deep ones take more time and are less likely to be detected. Moreover, Level 3 insights are more probable to generate an emotional reaction that combines frustration and excitement, while easily noticeable facts are less likely to be accompanied by excitement. Therefore, the probability of accurately capturing an insight by measuring the emotional response of a subject via EEG seems higher when the insight is deeper, more complex, and it occurs later in the analysis process.

As previously, the participants have been asked to complete a questionnaire after employing the visualizations. Questions that were posed involved the interaction and visualizations, as well as the relevance and accessibility of the data presented in them. Many participants suggested that they did not experience the Aha! effect. Reasons given for this included the fact that the information they discovered had a low complexity and was “easy to find”. Furthermore, while they were interested in the presented topic, they were not involved with it to the extent that any newly discovered insight would influence their way of thinking (“I don’t think this can surprise me”). When inquired about the moments of insight, participants mentioned that they reached some answers by a simple search process. As suggested by [11], logical search




Figure 4.15: Correlation between the insights generated by the participants and the emotional responses detected by the EEG headset. The insight moments are marked by 25 plus (+) signs and are plotted along the horizontal timeline. Note that all + signs are visible and there is no perfect overlap between two moments of insight. The colors of the insights represent the depth of the generated insight: blue is Level 1 (simple), orange is Level 2, and red is Level 3 (deepest). The green and red vertical lines beneath the insight + signs indicate the presence of an emotional response. A green line under a certain + sign indicates the presence of frustration previously to the insight generation. Similarly, the red line under a particular plus sign indicates the presence of excitement in the moment of insight generation. The three colored lines above the + signs represent the kernel density estimates for individual Gaussian kernels constructed around the three types of data points from the horizontal axis.

for new information is a process different for gathering knowledge than the one of epiphany. Based on the questionnaire results, an even stronger correlation was noticed between the instantaneous insights that would not involve a systematic search process and the emotional responses; but as these investigations were subjective—based on open-end questions and the verbal narration of the participants’ insights during the task—no quantitative values are currently available. These answers, together with the unpredictability of insights, could represent a partial explanation for the limited number of deep insights generated by the participants. Our hope is that further experiments can generate a larger set of insights in diverse visualizations, and thus offer a more complete view of the possibilities and limitations of mobile EEG detection of insight.

Another relevant aspect for the validation of EEG measurements for detecting insight moments is given by the false positives. In our scenario, false positives are represented by moments in time when no insight is generated, but the emotional states of frustration and excitement are detected inside the time-frame described in Section 4.2.2. In the second row of visualization tasks, only nine such instances were recorded by the EEG system. As the possibility exists that these were generated by elements external to the software system (real-world or mental distractions of the user), further investigation is required to deduce their impact on a larger scale. Note that an insight is implicitly considered true by the issuer, at least in the initial stage of the Aha! moment. Usually, in knowledge tasks a generated insight later undergoes a validation process that implies the systematic analysis of all facts and the insight information. This can result in an erosion of confidence, but even insights that contain false information will most likely have the potential to generate an emotional response. As a result, the EEG measurement of emotional states generated by insights should not be considered as a validation of the provided information, but as a sign for the presence of insight.

4.2.3 Discussion

One can hypothesize about the potential of EEG measurements—and in a wider sense of emotion detection—to accurately capture the presence of moments of epiphany or insight in subjects during various tasks, like problem solving, pattern recognition and extraction, concept manipulation, etc. Although the nature of insight and emotion is clearly subjective, the presence of a mechanism for connecting these elements and reliably detecting one of them through mobile brain-imaging [47, 49] opens an entire set of possible research directions for the future.

A major opportunity in this sense is represented by the quantitative detection of insights in the process of evaluating visual applications and information visualization techniques. Especially for information visualization methods, the capacity to generate insights is the essence of a powerful representation [201]. While emotional response does not quantify the power of an insight, it is capable of suggesting the presence of a reaction generated on insight. Additionally, this can suggest the relative value of the insight to the person, as our tests revealed that insights generate a detectable emotional reaction mostly if they are sufficiently deep, take a longer amount of time and effort to achieve and the topic of the problem is relevant to the subject.

Besides evaluation of visualization techniques, the capacity to detect the moment of insight can be used in collaboration (e.g, signaling how well participants are integrated in the process and managing the data flow) or for automatic operations (e.g., data tagging and binding based on the interactions the user executed shortly prior and during the moment of insight, highlighting of information involved in the Aha! moment and capturing layout screenshots that are relevant to a particular insight).

In the next two sections, we apply our findings on EEG measurements in order to visualize user affective states on two different types of interfaces: multi-touch and desktop. The proposed visualization solutions and their features are focused both on real-time representation (i.e., visualizing user emotional states during their interaction with the system) and post-task visualizations (i.e., representing user emotions after the interaction took place, for the purpose of analysis).

4.3 Emotion Visualization on Multi-touch Displays

Emotions are one of the omnipresent aspects of our lives: whatever we do, we are always accompanied, influenced and even defined by the emotional states that we experience. As such, it is not really surprising that scientists, technicians and developers are often seen as “cold” or rather emotionless, as informational systems are one of the few fields of our modern lives where feelings have gained minimal importance in the last years. This is partly due to the fact that IT systems are generally oriented towards information, speed and efficiency, but also because the interpretation of human emotions is as elusive and subtle as the emotional states themselves. Should this mean that software applications are doomed to contain only “dry”, quantitative information?

In the previous Sections 4.1 and 4.2 we have already started to highlight the importance of emotion recognition techniques, especially in the context of subjectivity measurements and system evaluations. Moreover, many techniques that can indicate the presence of certain emotions are slowly finding their way into software systems, most of the accent falling on affective applications [171] that are able to detect a set of user emotions and change system behavior based on these readings (i.e., emotional states as subconscious data input). Further examples for this include applications that react to emotions extracted from speech [160], as well as systems that change the style of a painting based on the emotional state of the viewer [261]. However, emotional feedback has a much broader potential that can be only fully achieved through the exploration of an additional path. More specifically, emotion recognition and classification techniques offer emotional feedback in order to enable affective awareness and self-regulation of both single user experiences and multi-user collaborative scenarios.

In the following, we propose a visualization approach, called *EmotionPrints*, for incorporating real-time information about user emotions in new or already designed multi-touch interfaces. Similarly to fingerprints, where touching an object would leave a marker on the object uniquely identifying the person, EmotionPrints aim at introducing a standardized visualization of emotions by marking the virtual object that a user touches with a representation of the user’s current emotional state. The power of this representation lies in the fact that it is applicable to any multi-touch screen and any interface element—regardless of its shape, size or color—as long as at least parts of its margins are inside the interface. Furthermore, our representation

can encode and display a variety of emotions in terms of affective valence (pleasant or unpleasant) and arousal (excited or calm), while at the same time being entirely independent of the technique employed for acquiring the user’s affective states. The goal of EmotionPrints is to improve the user’s emotional self-awareness, the awareness of other users’ emotions in collaborative and competitive scenarios—be it games, visualization systems or any collaborative touch-based application—and support the evaluation of touch systems and user experience through the visualization of emotional involvement.

In Section 4.3.1 we highlight the design considerations that an emotion visualization for multi-touch interfaces would need to observe. Then, we describe the details of our approach, both in the context of real-time visualization of user emotions and their post-task analysis. The results we obtained with EmotionPrints are presented in a user study, where user emotional states are acquired through wireless BCI headsets. Building on this, a discussion segment addresses benefits and limitations, as well as addresses potential concerns related to the topic of privacy.

4.3.1 Design Considerations

As a vital starting point, we have composed a list of requirements that a real-time visualization of user emotions on multi-touch interfaces would have to satisfy in order to offer an effective representation of affective states. The elements of the list are based both on related research and our previous research on the topic:

1. Less standardized than graphical user interfaces (GUIs) designed for desktop environments, graphical elements on a multi-touch device can have almost any shape. Accordingly, the representation has to be *generally applicable for any interface element* regardless of shape and size, and at the same time refrain from displacing or repositioning any of these interface objects.
2. The visualization of the emotions should be in the appropriate context, ideally *at the physical location of the touch event*. This would support easy identification of the operation-emotion couple as well as, in some cases, identify the corresponding user.

3. Human emotions are rarely precisely defined and complex in nature. A visualization that tries to convey emotional states or attributes should consider these inherent properties and offer representations for *hinting at the emotional state* of the user instead of trying to visualize exact values.
4. The representation needs to be as *intuitive* as possible and support *fast* processing. This is particularly important as users need to perceive and interpret the displayed emotions during the real-time touch-based interaction.
5. Various levels for emotional involvement have to be *distinguishable*. Also, users should have the possibility to compare the different emotions that are displayed at a point in time.
6. The representations should support the *real-time* nature of the interaction and also ensure that the interface does not succumb to visual clutter.

4.3.2 EmotionPrints

As highlighted previously, different sets of user emotions can nowadays be detected through a relatively wide range of techniques. The importance of detecting and classifying these emotions increases as one considers that user emotional states are usually generated by *stimulus events* [243], events which in our context can be related to the applications the users interact with. As a result, emotions might be closely connected to the visual systems users interact with (e.g., frustration about an unintuitive representation), to the task they have to execute (e.g., excitement after having an insight), or even to the other users they collaborate with (e.g., group tensions).

Therefore, we aim at creating an emotion representation that is independent of the emotion acquisition technique and is thus combinable with any approach that can interpret affective states in real-time. In order to demonstrate and evaluate our concept, our implementation and corresponding study employs emotion reports obtained through BCI readings. Details about the emotion acquisition process are highlighted

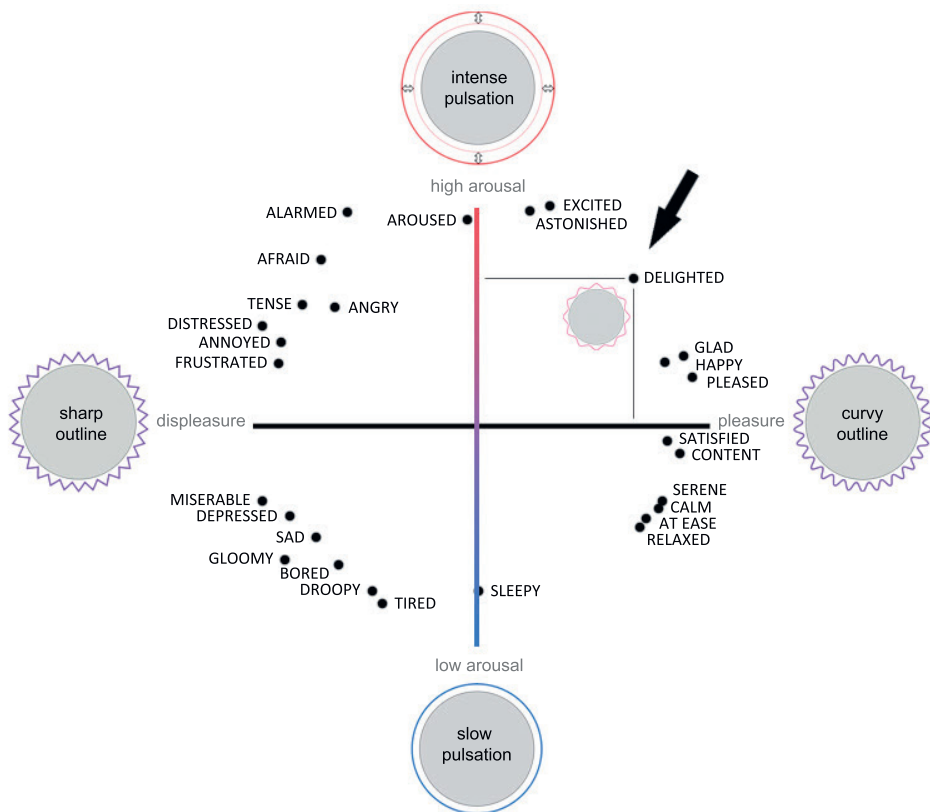


Figure 4.16: Russell’s circumplex model of affect (cf. Figure 2.1) extended by the visual metaphors employed by EmotionPrints. On the horizontal axis of valence, positive and negative emotions are encoded by outline texture. Positive emotions are represented by halos with an increasing number of wavy, round shapes, while negative ones have halos with jagged, angular outlines. The intensity of each emotion—positive or negative—increases with the number of shapes that define the outline of the halo, e.g., see arrow pointing at the state of *delight* and its corresponding emotion-print. On the vertical axis, high arousal is encoded by intense red color and a fast pulsation of the halo, while low arousal is encoded by a blue halo and almost no pulsation at all.

in the user study from Section 4.3.4. Furthermore, an initial prototype of the EmotionPrints design (detailed below) was implemented for the Multitouch Cell tabletop display. This prototype was coded in C++ with the help of the TUIO framework⁵ and Cornerstone SDK⁶.

⁵TUIO open framework defining a common protocol and API for tangible multitouch surfaces, <http://tuio.org> (February 2014).

⁶MultiTouch Cornerstone SDK, <https://cornerstone.multitouch.fi> (February 2014).

Once the emotional states are established, the design considerations could be addressed in order to generate the most fitting visualization. To satisfy the first two requirements, we decided to employ a metaphor based on the actual fingerprints people leave when interacting with touch displays. More precisely, similarly to [311], virtual objects that would be touched and acted upon by users would receive a colored halo that would closely follow their outline. Thus, the EmotionPrints would be included in the representation spatially right at the position where the touch event took place and without influencing the previous distribution of elements on the interface.

To model the halo effect around the objects we had to generate an offset-curve following the underlying shape (Figure 4.17). These halos would extend the outline of the virtual objects by a constant distance a , computed along the normal to the object's outline. At the object's edges, the halos would be rounded in order to offer a smoother contour capable of sustaining an additional layer of information: outline texture. With this approach, colored halos could be displayed for virtually any object on the display, regardless of shape and size. Note that the developers of multi-touch applications can control which of the virtual objects will be able to represent an emotion-print around their outline by adding the ID of those objects to a list of *touch_items*. Through this, developers have the option of defining a subset of virtual objects that can display the emotion halos, as not all interface elements might be relevant for enhancing user awareness or evaluating the interface and user experience.

Considering the inherent complexity and imprecision of emotions, we addressed the third design consideration by employing Russell's circumplex model of affect (Figure 4.16), as previously described in Section 2.2. Widely accepted in the literature of affective systems [96, 230], Russell's model allows us to take any emotion included in this domain and convert it to a valence-arousal pair (see [204] for considered values), which can function as the core data for our representation. Further, visually encoding valence and arousal instead of specific user emotions enables the use of a wider range of techniques for detecting user affective states. As long as these emotions would be part of the model, they could be visualized in terms of valence-arousal pairs.

The visual representation of EmotionPrints is partly inspired by the interface of the eMoto mobile application [267], that focuses on emotion expressivity and communication by using color, shapes and animation to encode the two-dimensional emotional space defined by Russell's model. Similarly to this approach, EmotionPrints employ




Figure 4.17: Computing the outline of the EmotionPrints halo: (left) The basic shape of the halo is computed by extending the outline of the touched object by a length a along the normal of the object shape; (center) In the case of convex edges, the edge of the halo is rounded, following the outline of a circle with the center in the edge of the object and a radius of a ; (right) For concave edges, the edge of the halo is again rounded by drawing a circle with the center X' and radius a . X' is computed as the point that is on the line defined by the object edge X and the virtual edge of the halo X_{EP} and positioned such that the following segments are equal: $\overline{X'X_{EP}} = \overline{XX_{EP}}$.

color, shapes and motion to represent various levels of valence and arousal (Figure 4.16). However, in contrast to eMoto, where the color attribute changes over the entire Russell model, our solution is intended as a simpler visualization, where users should be able to more easily and quickly distinguish the different levels of arousal and valence.

For the vertical axis of our representation model, arousal is double encoded through color and motion. As arousal increases from boredom towards excitement, the halos start to pulsate increasingly fast, in a manner that is meant to mimic the human heart rate in terms of frequency. At the same time, the color of the halos changes from blue to red as the perceived arousal levels increase. This is consistent also with Ryberg's theory [234], that suggests that emotional energy can be represented as increasing from blue to red. The two attributes of color and pulsation are meant to combine their effects on the human perception in order to transmit the metaphor of thawing and becoming excited or calming down and freezing, depending on the extremity of the axis.

On the other hand, the horizontal line of valence is encoded in the EmotionPrints representation through a variation of the halos' outline. More exactly, positive emotions are captured through an increasing number of waves forming on the outline of the halo, while negative emotions are displayed as halos with increasingly many jagged, sharp corners. Figure 4.16 highlights this representation of EmotionPrints for the valence-arousal space.

By employing different attributes for each of the two dimensions, as well as by using separable dimensions like color, curvature and motion [284, 304], we addressed the fourth design consideration by supporting the intuitive processing of the EmotionPrints. Furthermore, the attributes allow users to clearly distinguish between multiple levels of valence and arousal, thus also tackling the fifth design requirement. While it is clear that color and motion do not allow for a very fine-grained classification, our study has shown that users are capable of clearly distinguishing at least three distinct levels of arousal: blue for low arousal, violet for medium arousal and red for high arousal.




Figure 4.18: Example of EmotionPrints dissipating and being overwritten by newer touch instances. The first touch of the user generates a halo that suggests low arousal and a slightly positive valence. This emotion-print starts to dissipate over time by increasing the alpha channel of the halo. Before it can become completely transparent, the user touches the same object again and the old emotion-print gets overwritten. This time, the emotional readings suggest high levels of arousal and an increased positive valence. Finally, the new halo starts to dissipate again.

Finally, in order to address the last design consideration and support the real-time character of the visualization, all EmotionPrints slowly dissipate in a preset time interval after the touch event occurred (Figure 4.18). Additionally, one can set an upper limit for the number of fingerprints that can be displayed simultaneously. These two

features—dissipation and maximum number of displayed halos—ensure that EmotionPrints can be customized for the particular needs of any multi-touch application, such that the visual clutter introduced by this additional layer of affective information is minimal.

4.3.3 Post-task Analysis of EmotionPrints

As stated in the beginning of Section 4.3, one of the goals of our visualization is to support the evaluation of touch-enabled applications by offering visual feedback about the affective states of the interacting users. However, as previously described, EmotionPrints are aimed mainly as a real-time emotion visualization approach that requires the representations to dissipate shortly after each touch event of the user in order to reduce display clutter and maintain the visualization up-to-date. Thus, to further support evaluation and post-task analysis of user interactions and corresponding emotions, we extended the concept behind EmotionPrints to incorporate an additional visualization that allows users to inspect the temporal and spatial distribution of the touch-emotion couples.

After a session in which one or multiple users have interacted with an application on a touch interface, the EmotionPrints histogram can be displayed, as shown in Figure 4.19. The histogram offers a temporal overview of the touch events of a user and the corresponding valence-arousal pair for each such instance. To establish this correspondence, our current implementation is able to track the hands and arms of multiple users around the tabletop based on the infrared image captured by its camera and the orientation of the users' arms over the table. At the same time, a down-facing camera tracks the motion of the tabletop users around the display for the entire length of the session, ensuring the correct user identification based on their location and arm movement.

While the initial design conventions have been considered also for the histogram, there are a couple of differences in the way emotion is encoded. In order to better differentiate between emotional states with positive and negative valence but also due to the inherent design of histograms, the valence corresponding to each touch is encoded by the orientation and height of the vertical bars. The two sides of the histogram—positive and negative—are indicated by a circle and a triangle that make




Figure 4.19: Histogram representation of touch events and their associated arousal-valence values for two users. Users can inspect their own temporal distribution of EmotionPrints and compare it to other users by snapping together two views, as in this figure. The bars are positioned along the temporal axis, and their orientation and size encode the valence and its intensity. At the same time, the color of a bar encodes the arousal level experienced at that point in time, during that interaction.

the connection to the round and sharp edges of the EmotionPrints, as highlighted in Section 4.3.2. Moreover, touching one of these two glyphs will hide the corresponding half of the histogram, allowing a user, for example, to display and more easily compare only the touch instances that represented a positive valence. In terms of arousal, the color of each bar encodes the arousal levels for the touch instance, following the same color mapping from Figure 4.16. Again, the selected attributes for encoding the two-dimensional emotion data are some of the more separable ones [304].

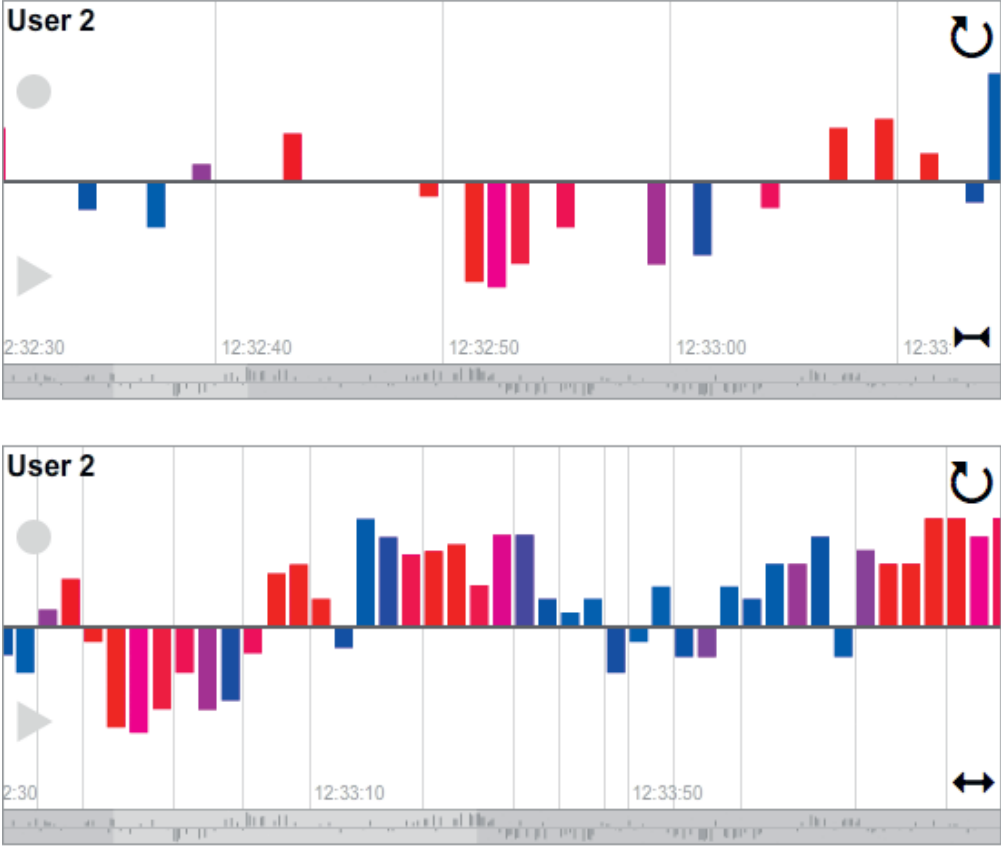


Figure 4.20: EmotionPrints histogram can be displayed with constant time intervals (top) or can be compressed to eliminate time intervals where the user did not execute any touch events (bottom). Switching between these two representations is done through the double-ended arrow at the right-bottom corner of the histogram. The curved arrow at the right-top corner allows users to resize and reposition the view.

As sometimes multiple touch events can be executed by a user in a very short period of time (e.g., fast tapping in a game), we decided to discretize the temporal axis and convolute all emotional readings for time intervals smaller than a second. For example, if a user executed three touch events in a one second interval, the valence and arousal values for these three instances would be averaged and only one bar would be displayed in the histogram (see Figure 4.21 where one histogram bar corresponds to two touch events). This allows us to reduce clutter, especially as emotional signals do not fluctuate with such a high frequency.

In terms of interaction, the histograms have a custom scrollbar attached, that allows users to both move to a certain location in the history and to see an overview of a user's valence during an entire session. As presented in Figure 4.19, two histograms can also be coupled in order to facilitate comparison. More precisely, when a histogram is positioned next to another one, the one that is being moved resizes to the same dimension as its counterpart and snaps into place next to it. Additionally, if there are common periods of interaction for the two histograms, the time axes align automatically.

Sometimes the density of touch events fluctuates heavily during a session. To allow for a more comprehensive exploration of tendencies and patterns, EmotionPrints histograms can be compressed to discard the temporal segments where the user did not interact with the interface (Figure 4.20, bottom). To maintain a frame of reference, the vertical interval lines are shifted such that areas where the lines are closer together suggest periods of time when the touch events were more sparse, and vice versa.

While the histogram encodes the temporal distribution of touch-emotion events for a selected user, it gives no information about what areas of the display or what objects have been interacted with. In order to achieve this, the individual bars of a histogram can be selected, highlighting the spacial distribution of the corresponding touch instances (Figure 4.21). This means that for each valence-arousal pair displayed in a histogram, one can inspect the ID of the issuing user, the location on the screen where the touch event/s took place and the operation that the user executed (e.g., touch, drag, resize).

Further, to obtain an overview of the spatial distribution of EmotionPrints during an entire session, a heatmap can be displayed over the area of the application. Again, this heatmap employs the same color mapping from Figure 4.16, and combined with the histogram visualization offers a comprehensive solution for user post-task analysis and evaluation support. However, one drawback of the heatmap representation is that it can only use colors to convey information. As such, one can only show the spatial distribution of either arousal or valence at a certain moment.

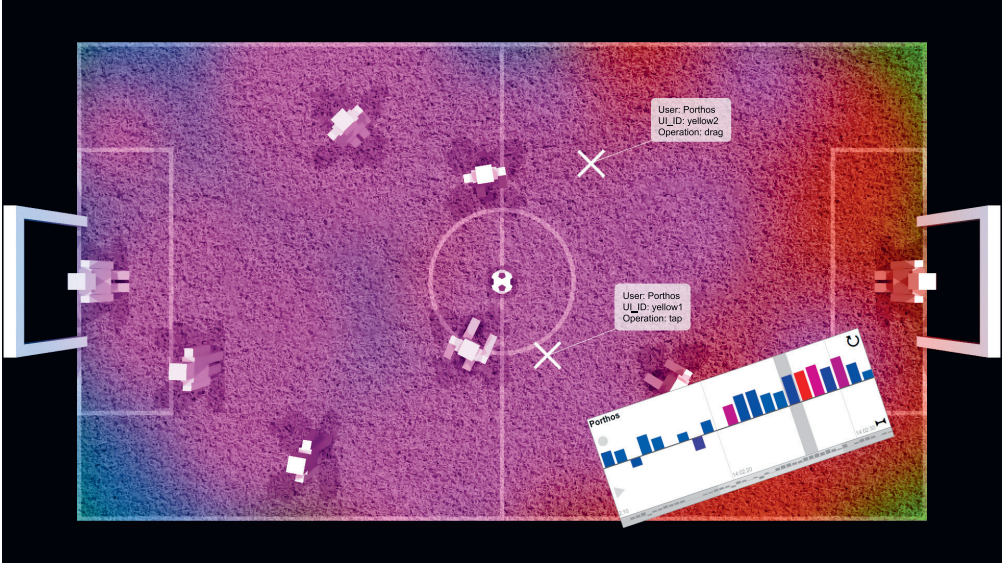


Figure 4.21: Screenshot of the eSoccer multi-touch game supporting up to four users. The background of the field is covered by a blue-violet-red heatmap that encodes the arousal level of a user during the last gaming session. An EmotionPrints histogram is displayed for the same user. By selecting a bar in the histogram, the system displays the quasi-simultaneous locations where the user has executed the corresponding touch operations on the screen. Besides location information, the histogram instance further correlates with the user ID, the ID of the touched object and the executed operation (e.g., tap, drag or resize).

4.3.4 User Study

When working with concepts like affective states, it is difficult to devise a study that can quantitatively inspect the merits of a visualization. In such a context, the issue of “correct” and efficient representation can be defined only vaguely, as it seems more important to evaluate whether the EmotionPrints were perceived by the users as intended in the design considerations, and if the representation had an impact in terms of emotion awareness and user experience analysis.

In order to validate our visualization, we incorporated EmotionPrints in two tabletop applications: a soccer game supporting up to four players (Figure 4.22) and a multi-touch enabled version of WebComets (Figure 4.23, see Chapter 3). In the soccer game, the only objects that users could manipulate were virtual soccer players, which were




Figure 4.22: Two players interacting with the eSoccer game on a tabletop while their emotional cues are being interpreted via BCI and represented through EmotionPrints, all in real-time.

then registered as *touch_items*. Similarly, in the touch-version of WebComets, the nodes of the graph were tagged as *touch_items* on which EmotionPrints could be displayed. The two applications were selected in order to investigate how EmotionPrints are perceived in a competitive and a collaborative scenario. Note that the alterations to the visual representation of the multi-touch version of WebComets are highlighted in Section 5.2 in detail (e.g., rectangular representations, multiple history windows, etc.), and that the interactive capabilities of the touch-based visualization have only been extended.

One difference between the two applications in terms of emotion visualization was that for the collaborative scenario, touches corresponding to a negative valence would not be displayed. The motivation for this is given by the fact that reinforcement and feedback on positive emotions has been shown to foster the collaboration process as well as contribute to the development of positive group affective tone [94,95]. Note that the group affective tone is defined as a consistent emotional state throughout the members of a group. More importantly, the presence of a (positive) group affective

tone can also influence the effectiveness of said group. This has also been addressed by Saari et al. [236], who state that in a group, negative emotions can lead to withdrawal, while positive emotions can foster the interaction and collaboration of the team members.




Figure 4.23: User interacting with the browser history visualization on a tabletop. EmotionPrints are being displayed each time a user touches a rectangular segment, corresponding to the nodes of the history that represent the websites that have been accessed.

In both applications, whenever users would interact with a virtual object, their current levels for arousal and valence would be determined based on brain signals obtained through EEG readings, transmitted to the tabletop computer and represented through an instance of EmotionPrints around the corresponding UI object (Figure 4.24). In order to detect the two dimensions of the Russell's emotional space, we again employed the EPOC device. In our study, coupling both the facial expression readings and the classification of affective states generated by the EPOC framework supplied us with values for the emotions that are closest to the extremes of the arousal-valence axes in Russell's circumplex model, and which can be employed to cover the entire corresponding 2D space.




Figure 4.24: User playing on a tabletop while wearing the BCI headset (left). Each time the user touches the multi-touch surface, the EEG signals are read and interpreted as affective states. Subsequently, EmotionPrints are displayed on the tabletop around the object that the user touched. The EmotionPrints dissipate in time in order to reduce display clutter (right, top to bottom).

Twelve subjects took part in our study, all of which had some level of experience with multi-touch devices or tabletops. The participants were split into six pairs, where each pair was asked to play a game of virtual soccer against each other and to collaboratively find a node with a particular set of attribute values in the history visualization. Before the tasks commenced, each team got introduced to the functionality of the

two applications until they felt confident to use them. All participants were equipped with an Emotiv EPOC headset, but only after the abilities and limitations of this system has been highlighted to them, and they agreed to share the readings and interpretations of the device.




Figure 4.25: Computational path for the EmotionPrints representations: (left) Gathering data that can encode information about user affective states (e.g., EEG signals, facial expressions, galvanic skin response); (center) Interpreting this data as affective states that can be decomposed into valence and arousal; (right) Representing the valence and arousal information on the multi-touch surface as EmotionPrints.

In order to compare the impact of EmotionPrints, the six groups were divided in two, where three groups would complete the two tasks enhanced by the emotion representation obtained through the users BCI readings, and three groups that could employ the same two applications without EmotionPrints representations. However, the members of the groups that did not have the benefit of emotion representations did still wear the BCI headsets. This allowed us to inspect and compare user behavior in cases where similar valence-arousal readings would be reported. The members of the pairs that would work with EmotionPrints were given a thorough presentation of its design and functionality. Furthermore, all users were encouraged to express themselves freely during the tasks.

During the study, almost all participants engaged in verbal communication, allowing us to analyze their interaction relatively to the EmotionPrints. In the context of the soccer game, there were multiple instances where the groups that had the emotion visualization at their disposal managed to use this information, in some cases even gaining an advantage. For example, a player that was leading 4 to 3 after being down 1 to 3 saw that his opponent's EmotionPrints were very edgy and red, thus making the

following statement: “Wow, you must be really annoyed”. He later was able to extend his lead also by directly addressing the frustration of his opponent. On the other side, in the collaborative history visualization, users that had access to viewing the EmotionPrints would often notice the positive excitement of their colleague, suggested by phrases like “Did you find something?” and “You think it’s in this area?”. In contrast, in the pairs that did not have access to the emotion representation, the awareness of affective attributes was minimal. This is also supported by the fact that although users in these three groups experienced similar emotional configurations, there were almost no discussions addressing real-time user experiences.

After the two tasks, the participants that worked with EmotionPrints were asked to fill out a questionnaire about their design and utility. The subjects considered that the design was intuitive and distinguishable. However, some would have liked to adjust the time interval in which the prints would dissipate. All participants considered the opportunity of enriching the interface of multi-touch systems with emotional states a good idea. The most powerful feedback was related to the impact on collaborative work and competitive gaming on multi-touch devices (“I really want to see this in some of my iPad games”). At the same time, the collaborative aspect brought up issues like privacy and property of the represented emotions, which are further addressed in the following discussion segment.

Furthermore, all 12 participants had the opportunity to inspect and analyze their emotional readings in the EmotionPrints histograms and heatmaps. Their verbal iterations during the analysis process suggested that most participants managed to gather valuable insights about their emotional states, increasing the overall awareness of their own and their colleagues experience (e.g., “I really got frustrated after that, didn’t I?”, “What did you see there?”). At the same time, the heatmaps gave them insights both about the areas that they most interaction one and their emotional experience during those touch events. For example, one user preferred executing his attacks on the side of the field in the soccer game, and was mostly successful in this approach. Thus, the heatmap for this case presented more intense coverage and positive valence in those areas of the field.

4.3.5 Discussion

As briefly mentioned in Section 4.3.4, the participants of our study have addressed the issue of emotional *privacy*. On one hand, the hardware/software solution has to ensure a basic level of privacy and security. This topic has been studied in [179], where data leakage from EEG devices has been used to gain information about private user information, credit cards and PIN numbers. There are currently no known studies on the security of EEG devices in the context of emotional readings. On the specific topic of the Emotiv BCI, the data obtained by the headset is encrypted for the communication with the corresponding computer, thus making it less susceptible to third party attacks.

On the other hand, perhaps similarly to social media and other approaches to digital sharing, emotion visualization both in single-user and collaborative context has to commence by informing and empowering the user. Irrespective of how user emotions are interpreted and represented, users need to first be informed about the particularities and capabilities of the acquisition technique, be it emotion recognition through facial expressions, BCI, or physiological responses. Only through this, users will be able to take a pre-task informed decision whether they want to share this data and to what extend.

This is the reason why applications that are enhanced with EmotionPrints should inform the user about the interpreted emotional states and require them to agree with those terms, as in any virtual sharing activity that involves private data. More precisely, users should be made aware that the visualization approach does not inform other participants about concrete affective states, but rather converts this information to the space of valence and arousal. Furthermore, as in the collaborative history visualization, Russell's two-dimensional space can be further restricted to filter out a set of values.

Note that the visual representations introduced by EmotionPrints are consistent for any emotion detection or signal interpretation technique that outputs either directly user valence and arousal levels, or reports the presence of one or more emotions included in Russell's circumplex model of affect and that can be decomposed to the two-dimensional encoding.

Other concerns that have been raised include display clutter and a corresponding increased cognitive workload. To address this, the last design consideration has been added to our concept. Limiting the period of time that an emotion-print stays visible has multiple advantages: the presented emotional information is never outdated, the display objects can quickly revert to their original form in order to minimize clutter, and the identity of the user that generated the print can be better recalled compared to an emotion-print that has been displayed for a longer period of time.

4.4 Emotion Visualization on Desktop Interfaces

Chances are that almost every person who can claim to have worked with a computer in recent years would immediately recognize the graphical user interface (GUI) of a desktop environment. But why is this the case? How come this paradigm of a pointer, of windows, of text fields and buttons has gained such deep roots in our minds? For sure, one of the reasons for this is the standardization of the desktop user interface. Today, regardless of what operating system or development environment users employ, they have access to the standard UI components, all with roughly the same properties and functionality. But while this uniformity has its advantages, it also has its pitfalls; one of which is the difficulty of combining these standard components with a higher level of flexibility. At the same time, this rigidity of desktop GUIs means that they tend to remain solely means for interacting and communicating purely objective information and navigational commands between the user and the computer.

In recent years, researchers have considered enhancing the GUI with subjectivity and emotions [114, 162, 180, 189] in order to enrich the user experience and support analysis, collaboration and awareness. However, most of the proposed solutions aimed at modifying the design of the common desktop by adding virtual (emotional or social) agents or by offering a highly dynamic interface capable of changing based on the user's emotions. In the following, we propose an approach aimed at enhancing the standard UI components with minimal emotional representations. Called EmotionScents, these emotion-enhancements for common GUI widgets can encode and visualize the emotional states experienced by users while interacting with these wid-

gets. Similarly to EmotionPrints, we believe that changes in system properties can have effects on the emotional states and reactions that the user experiences [280]. As such, our hope is to capture user emotions in real-time during the desktop interaction and offer the user a per-component overview of the experienced subjective states.

The motivation of this research consists in the fact that an additional level of emotional information included in the GUI can, as already shown for EmotionPrints, support the user's emotion awareness as well as, similarly to information scents [56], improve post-task analysis and decision making. Building on the concept of EmotionPrints that is aimed only at co-located collaboration, EmotionScents has the potential of offering a viable alternative for emotional awareness of users engaged in distributed communication and collaboration, similarly to emoticons. Moreover, the focus of the proposed visualization framework considers two aspects: enriching the GUI with user emotions and visualizing these emotions in correlation with the most probable event (or interface component) that generated or influenced them. Such a combination can guide the user towards new assumptions and discoveries that otherwise would not have been explored.

In terms of utility, EmotionScents are aimed at offering support for real-time emotion awareness and post-task decision making and analysis. The visualized emotions can also have different sources: these can be the emotional states of the current user that also employs the scented interface, or the emotions of other users recorded in previous sessions. In the first case, the goal of the scents is to offer hints to the users about their emotions and behavior. This way, users can have a better awareness of their emotions during the interaction with a UI component. Further, this increased awareness coupled with the correlation between the executed event and the recorded emotion can suggest links between emotional states and user decisions (e.g., selected one radio button or another) or reactions (e.g., clicking a button changes the amount or filtering of some information enabling the user to gather a novel insight) [280]. In other instances, it might be useful to visualize a temporal distribution of user emotions for the same application or different versions of the same application, offering users an overview of their past experiences and emotional states when interacting with the various components of the GUI.

But emotional self-awareness is not the only target for EmotionScents. As suggested above, scented GUI widgets could also incorporate the emotions recorded in previous sessions by other users. In such cases, the experiences of a community of users can be detected and harvested, allowing following users to perceive the collective emotions and reactions of previous users that had to interact with the same UI component (i.e., make the same decisions, obtain the same data). For example, a user would need to manipulate a slider that is enhanced by the emotions of other users recorded in previous sessions. By means of EmotionScents, the current user can now see that most people selected values towards the right end of the slider (90-100%), but at the same time, most of them experienced negative emotions while doing so. On the other hand, those that selected slider values around 20-30% experienced more pleasant emotional states. This information, correlated with the information about the application type and the concrete information that the slider is manipulating, should offer the current user the ability to make a more informed decision while interacting with the application.

The fact that in certain cases EmotionScents should represent previously gathered emotional states means a decoupling of the emotion acquisition process from the visualization through scents. More precisely, in instances where the visualized EmotionScents represent real-time emotional states the user needs to be subjected to certain emotion detection methods (in our case BCI-based detection), while in cases when the user simply executes a post-task analysis or wants to use the collective emotional experiences of previous users for decision support, this is achieved by accessing the emotional database containing the previously recorded emotion sessions.

In the following subsections, we first highlight the design considerations that EmotionScents need to satisfy, followed by information on the recording and storing of user emotion data as well as a description of the visual representation. Further, we highlight the functionality and possible advantages of EmotionScents through a use case and an evaluation. We conclude with a discussion on current limitations.

4.4.1 Design Considerations

Before a GUI enhancement like EmotionScents can be implemented, one has to consider the effects it can have on the already existing interfaces and the methods that can be employed in order to transmit the emotional information to the user. Thus, a set of design requirements should be established that consider particular aspects of the problem: the type of emotional information metrics, their visual encoding and their inclusion in the (potentially preexisting) interface. The following design guidelines have been derived through a set of preliminary studies, related research and the personal experience of experts in the field of GUI design, user experience and human-computer interaction.

1. The desktop GUI should suffer *minimal changes* in the process of adaptation. More precisely, if an application interface already exists, the EmotionScents should be implemented on it without disturbing the layout of the window or changing the positions and sizes of the UI components.
2. For the same type of GUI widgets or widgets that are similar in either functionality or representation, the EmotionScents need to be located in the *same relative location* to the widget. This aids in making the scents more identifiable, while at the same time allowing users to easily establish the correspondence between scent and UI component.
3. Together with the positioning, the size and shape of the scents should allow for *comparability* between the various instances in the interface. Thus, when inspecting emotion-enhanced UI components, the user should be able to differentiate, analyze and compare their content.
4. The user should have the possibility to activate/deactivate the visualization as well as *customize the emotional representation* by changing data dimension priorities, employing temporal filtering or selecting a custom range of users for which to represent previous emotional states.

5. Due to the restrictions in available GUI area and the necessity to maintain the initial aspect of the interface, EmotionScents should offer a *cumulative overview* of the emotions experienced by the user or users for the corresponding UI components. As, for example, clicking a button can happen hundreds of times in a session, or even thousands of times in a set of user sessions, it seems more important to offer a high-level view of the experienced emotions than to include clearly browsable emotion information for every event instance.
6. The EmotionScents representations should be able to clearly and concisely *reflect the various dimensions of emotional attributes* and their distribution over the user space. In other words, while it is not the goal of EmotionScents to give the ability to inspect details about every pair of event-emotion instance in the GUI, they should still allow users to differentiate between various facets of the experienced emotions (e.g., pleasure and frustration versus excitement and calmness).
7. Starting from the previous assessment, one should *avoid considering too many emotional dimensions* at the same time. As available space and the shape of the enhancements are vital to avoiding changes in the interface, finding the right amount of emotional and statistical information to be displayed is of key importance.

4.4.2 EmotionScents

Before highlighting the details of the visualization framework, we need to again address the emotion acquisition process. Similarly to EmotionPrints (Section ref:emotionprints), EmotionScents employ the Emotiv EPOC headset in order to obtain real-time affective information about the various users. The reported emotions from the EPOC framework are used as direct input for Russell’s circumplex model (see Figure 2.1) in order to decompose them in terms of valence-arousal dyads. At the same time, using a model like this one that maps a large set of common emotions to a 2D space allows us to again decouple the visual representation of emotions on GUI elements from the medium that executes the actual detection. This means that EmotionScents could be used together with other techniques for emotion detection, like real-time interpretation of facial expressions, electrodermal activity or voice analysis.




Figure 4.26: Emotion data acquisition and representation loop: (1) Whenever the user interacts with an input GUI widget, the system records the EEG brain signals via the EPOC headset; (2) The signals are then interpreted as emotional states and stored with a reference to the generating GUI widget; (3) The emotions captured this way can be visualized on their corresponding widget via EmotionScents, in real-time for emotion awareness or post-task for analysis and decision support.

From a procedural standpoint, Figure 4.26 highlights the acquisition and visualization loop of EmotionScents. In a first step, the user is interacting with the GUI widgets of an application or system. In parallel, the EPOC device is being deployed and continuously reads the user’s brain signals. If the user interacts with a particular UI component (e.g., a button) at that instance a signal is sent to read the current brain signals; these signals are then interpreted as emotions by the framework on the computer and the results are stored in a database. Depending on the selected mode, these emotions can be then represented in real-time on the UI component the user interacted with. Alternatively, users have the possibility of querying this database of collected emotions and show EmotionScents for entire sets of emotional experiences.

Furthermore, there are different types of emotional triggers that have to be considered: “triggers that result in an action (physical response) or a reaction (emotional response)” [80]. More precisely, the assumption we make is that both emotions that the user experiences slightly before and slightly after the interaction can be related to the action he/she exerted. This is due to the fact that users might experience affective states on action (e.g., because of a previous decision) or on reaction (e.g., due to new information that has been added to the system). As such, emotional states are recorded for a couple of seconds before, during and after the interaction. Later, the user that displays this information has the possibility to customize the EmotionScents by selecting between the modes *emotion on action* and *emotion on reaction*.

User emotions are not the only information that is being stored in the emotion dataset of an application. Besides the valence and arousal values that are detected when a user interacts with a UI component, the system also stores the exact time when the interaction took place and the identifier of the UI component for subsequent correlation. Finally, in the case of UI components where the user can select multiple states (e.g., slider with multiple values, combo boxes, lists, etc.) an additional value is stored representing the selected value. On one hand, this is relevant because the emotion of the user should be coupled to the actual event and decision; on the other hand, as highlighted in the following paragraphs, GUI widgets can have separate *emotion scent* representations for each available component state.

Guided by the considerations from Section 4.4.1, the EmotionScents (Figure 4.27) were designed to encode two emotional dimensions (valence and arousal) as well as for certain UI components the user selected values. Their layout in the GUI is similar to the one presented in [14], while their composition resembles to the visualization of Twitter sentiments in [303]. Scents are represented as thin, horizontal, colored bars, composed of horizontal segments of equal length x , see Figure 4.28. The overall width n of the bar is constant over all the EmotionScents in an application in order to allow better comparability. As such, the size of a segment x is computed by querying the database for the number of distinct emotional readings available for the current control. For example, if an emotion bar has the width of $n = 60$ and there are ten distinct emotional readings in the dataset for the current control and timeframe, the width of a segment will be set to $x = 6$. This also means that the higher the number of available emotion samples for a scent, the smoother seems the transition between the colors as segments can take widths of one pixel or less (Figure 4.27c).

As each bar segment represents an actual emotional reading, the encoding of the nuance of the emotion (positive or negative) or the overall excitement is done through colors. But a particular difference between the two emotional channels of valence and arousal consists in the fact that while valence is perceived to be bipolar, arousal on it's calm-excited scale seems to be a more unified characteristic, as suggested by our initial discussions with GUI designers and HCI experts [319]. Thus, we encoded the two axes as following: red for negative valence; green for positive valence; and blue for excitement. Similar settings have been used in emotion-related applications [122,202], as using the trichromat characteristic of the human eye should improve the visibility of the colored areas—a vital aspect for visualizations with limited display space.




Figure 4.27: Different representations of EmotionScents: Two buttons with their corresponding emotion scent representations for valence (top-left) and arousal (bottom-left). Two check boxes with the valence and arousal EmotionScents displayed in parallel. Depending on the sort priority (i.e., based on valence intensity (top-center) or arousal intensity (bottom-center)), the position of the segments in the second dimension are established by their correspondence to the sorted ones. Two radio buttons presenting their EmotionScents encoding valence (right). The top bar includes over 50 separate EEG readings, thus resulting in over 50 segments that—depending on the interface—can have a width of under one pixel, resulting in a smooth transition between colors and suggesting a high number of emotional readings for that particular widget. For the bottom bar the segments are more visible, as they encode only 10 EEG readings, resulting in wider segments (i.e., larger x values).




Figure 4.28: The emotion scent representation. The length n of the bar is predefined. Each colored segment encodes an emotional state measured during the user’s interaction with the current GUI widget. Green segments represent positive emotions, while red ones encode negative emotions. In this particular representation the segments are sorted based on the valence of the emotions in order to allow a better recognition of the emotional distribution. Note that the length of one segment x is inversely proportional to the number of emotional states stored for the current widget ($x = n/n_{readings}$).

Further, scents are positioned immediately over the lower-left corner of the corresponding UI component (Figure 4.27). However, the user has the possibility of adjusting both the size and positioning of the scent to fit the layout and needs of the application. This is achieved through a context menu that the user can open by right-clicking on any *emotion scent* in the application.

Changes executed in the context menu of an *emotion scent* can have both local (only this scent instance) and global effects (all EmotionScents in the application). One important attribute that can be modified is the length of the scents, see Figure 4.29. The bars are set by default to equal lengths, supporting the comparison of ratios between positive, negative and aroused states. Still, if users are interested in focusing on the ratio of existing emotional readings for each UI component, they can switch to a view that computes a constant segment width x and variable bar with n .

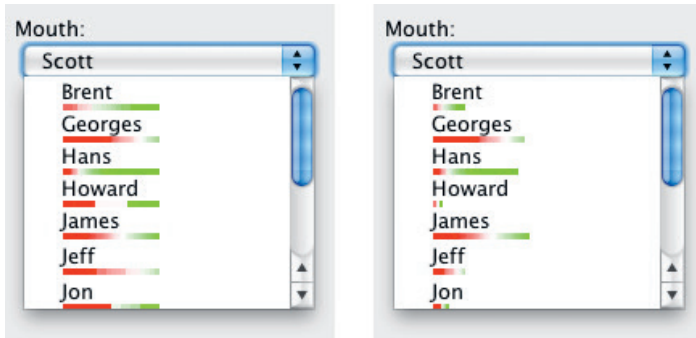


Figure 4.29: EmotionScents representation applied for a combo box control: (left) Normalized representation where the width of each bar is equal to the maximum width n , independently of the number of emotional readings available for each combo box item (i.e., n is constant). Such a representation enables a better ratio comparison; (right) Relative representation where each emotion segment x has a constant width computed as $x = n / \max(n^i_{readings})$. This representation supports the comparison of the emotional states quantified by the ratio of users that interacted with the corresponding items.

Moreover, the user has the ability to choose between representing either one emotional dimension (Figure 4.27 left) or both valence and arousal in the same bar (Figure 4.27 center). In both situations, the context menu allows emotion segments to be sorted based on different criteria: based on the intensity of the emotional channel (Figure 4.27) or on the temporal order of detection (Figure 4.32). A special case appears when both valence and arousal are displayed simultaneously. As the two emotional dimensions recorded for a particular user are closely coupled, one should consider this connection when representing valence and arousal next to each other. Therefore, in settings where both axes are visualized, the user has the possibility to sort only based on one emotional intensity, while the ordering of the second dimension results from the correlations with the first. Figure 4.27 shows two check boxes enhanced

with EmotionScents. In the top scent, for example, the valence bar is sorted by the intensity of the emotion (starting from very negative to very positive). The arousal segments from the blue bar below are sorted simply by establishing the correspondence to the same emotional instance/reading. This approach can have certain advantages, as the *emotion scent* for the top check box shows: while there seems to be a slightly higher prevalence of negative emotions, the users that experienced positive feelings while interacting with the widget have also experienced higher levels of excitement.




Figure 4.30: EmotionScents represented for slider widgets. The particularity of these widgets consists in the possibility of having multiple different values for a single GUI item. The top two sliders support the representation of EmotionScents by displaying a single colored bar that follows the slider head and encodes the cumulated emotions for the corresponding slider values. In cases like the bottom slider, where the distance between the possible slider head positions is large enough to accommodate the width of EmotionScents, multiple bars can be displayed at the same time.

As previously mentioned, for some UI components the user has the possibility to select a value. One such component that we considered is the slider (Figure 4.30). In order to maintain consistency and avoid visual overload, for sliders only one single emotion scent is represented that follows the position of the slider head and indicates the emotional states of the users that made the same selection. In cases with sliders that allow selecting a high number of values (40+), the scent can be customized to consider a window region around the current selection of the slider head and display all the emotions for users that selected values inside this window. On the other hand, if a slider has very few selectable values (Figure 4.30-bottom), EmotionScents can be displayed for the entire length of the slider, encoding the emotional states of the users that selected values in the corresponding selection windows.

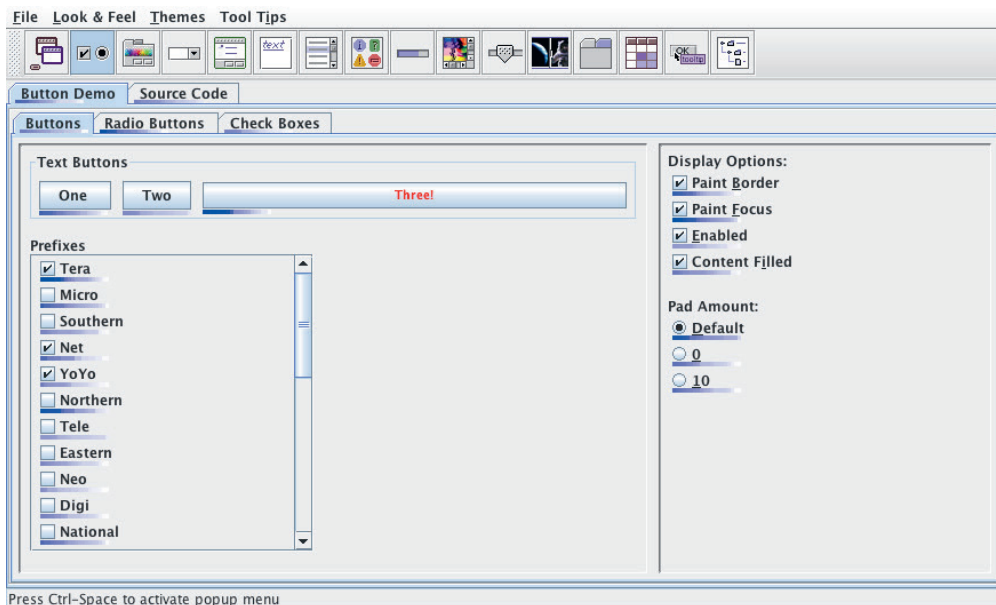


Figure 4.31: EmotionScents encoding the user arousal during the interaction with the SwingSet demo interface⁸. For certain GUI widgets, one can detect a greater incidence of higher arousal values.

In terms of implementation, the EmotionScents have been implemented in Java using the Swing GUI⁹ widget toolkit. The representation was obtained by implementing an alternative Look&Feel (or theme) that incorporates these colored scents. After the *emotion scent* theme is included in an application, the user has the possibility of enabling this interface by the simple click of a menu button. In the current version, only input or navigational widgets have been enhanced with emotion bars (Figure 4.31): buttons, combo boxes, check boxes, radio buttons, tabs, lists, tree views, sliders, and spinners. This is partly due to the fact that the user does not need to interact with all GUI widgets (e.g., labels or icons), which in turn means that in some cases there is no possibility to establish a correlation between an event and an emotion.

⁸<http://java.sun.com/products/plugin/1.4/demos/plugin/jfc/SwingSet2/SwingSet2.html>, Java SwingSet demo interface (April 2012).

⁹<http://www.oracle.com/technetwork/java/architecture-142923.html>, A Swing Architecture Overview (February 2014).

4.4.3 Use Case

Figure 4.32 highlights a use case for the EmotionScents visualization: A programmer called John starts a coding session in his Java development environment. During a two hour period in which he is wearing a mobile BCI and/or is recorded by a facial expression system, he extends an already existing code by making changes in multiple code files, and compiles the application several times. At the end of his session, he goes to the *view* section in the menu bar and selects another Look&Feel (i.e., theme) that displays the emotion cues inside the application GUI (Figure 4.32). By right-clicking on any of the EmotionScents, he can filter the time period for which the emotion data is considered and customize their representation through a context menu. After selecting the data for the last two hours, he notices based on the valence representation that he accessed two files very often and that in most cases his experiences when selecting these files were positive. At the same time, John observes that the debug button was executed quite often and that especially in the second half of the session, debugging operations were becoming frustrating to him. A similar pattern is valid for the execution button. He thinks back and remembers that he had to handle a set of errors that appeared due to his initial design. John switches to the arousal visualization and notices the same pattern: higher excitement value for the second half of his session in the case of the debug and run buttons. He now realizes that the errors he was trying to solve put him under more stress than he initially realized. Thus, he decides that in the future he should invest a bit more time in the design process to avoid unnecessary stress and obtain a better productivity.

A similar example, where the emotion information of *multiple* users is visualized, can be given with the following case. A company is creating a complex software solution. The developers and designers want to find out which modules and corresponding interfaces of the application induce frustration to the users. For this purpose, during the testing stage they collect—with the consent of the users—anonymous emotion data correlated with the various GUI widgets present in all the windows of the application. With this emotion database, they can now inspect what events produced intense emotions on user-side and analyze how these states might be influenced by user decisions, system interaction or displayed information.

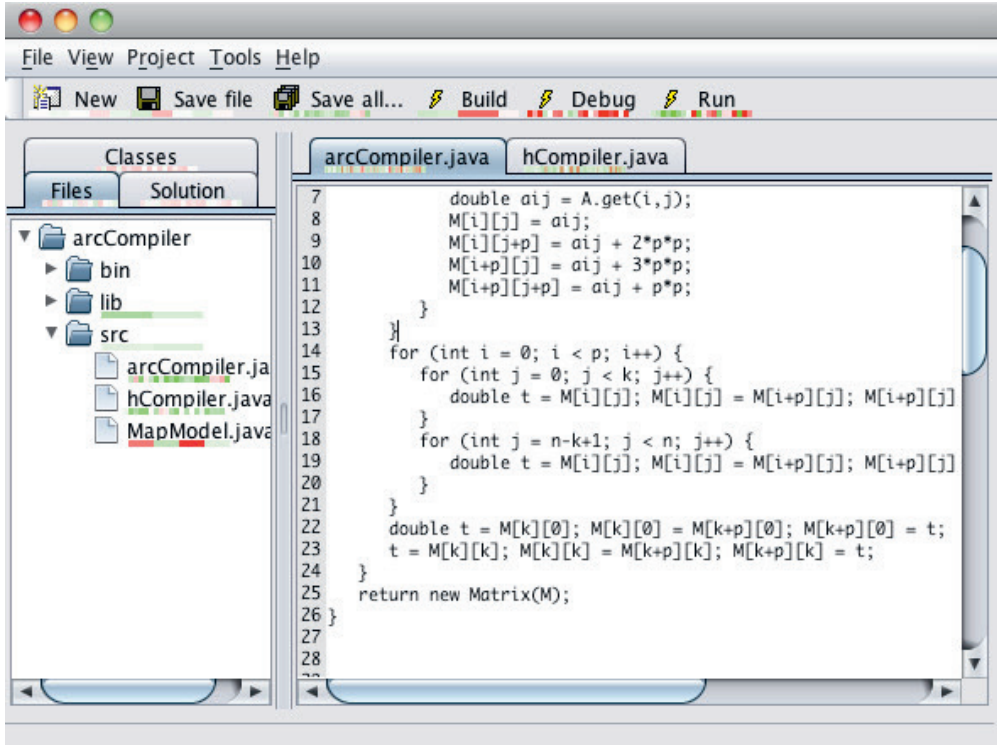


Figure 4.32: EmotionScents displayed on the interface of a simple Java IDE. The colored bars represent the valence of the emotion and the emotional states are in temporal order. The multiple red segments on the debug button could suggest that the programmer experienced difficulties in his coding process, especially in the second part of his coding session.

Certainly, in both examples the EmotionScents visualizations work only under the assumption that a method has been employed in order to record the emotional states of the user. While emotion detection based on facial expressions or speech is also a viable alternative, we focus on mobile BCI technologies in order to simplify the emotion acquisition process throughout this thesis. Even if current BCI approaches can be still cumbersome by requiring wet sensors and exact positioning on the scalp, EmotionScents is aimed partly at future generations of these mobile headsets, where using a BCI could already become as simple as putting on a hat.

4.4.4 User Study

An evaluation was devised in order to capture some of the advantages that an emotion-enhanced interface can offer. For this purpose, a tool was implemented that visualizes the U.S. foreign aid over the last decades, see Figure 4.33. The visualization was implemented in Java, while the interface was coded in the Swing widget toolkit; the represented data was obtained from the ManyEyes¹⁰ website. Additionally, the standard GUI widgets have been extended to implement EmotionScents.




Figure 4.33: Visualization for a dataset obtained from the ManyEyes website showing the distribution of U.S. Foreign Aid in the recent decades. The GUI widgets have been enhanced with EmotionScents to give another dimension of information to the navigation and analysis process of the data.

The evaluation included a group of 24 participants (9 women, 15 men) with ages between 19-55 years and an average of 28.25 years. The group consisted of people with at least intermediate experience with computers and mostly with previous experience in working with visualizations. After a brief introduction to the visualization tool, the group was randomly subdivided in two equal sets of 12 participants. The members of both groups received the task of gathering insights from the proposed visualization. More specifically, they were asked to inspect if they could detect any patterns between armed conflicts that the U.S. were involved in over a period of 50 years and the distribution of foreign aids. One of the simplest possible connection that we were looking for was the detection of an increased foreign aid to the nations with which the U.S. has been previously involved in a conflict.

¹⁰ManyEyes, <http://www-958.ibm.com/software/data/cognos/manyeyes/> (February 2014).

The members of the first group were asked to undertake the task first. Further, these participants were equipped with an EEG headset in order to record their emotional states during their interaction with the interface. However, none of them did benefit from having any EmotionScents displayed on the interface. Also, participants were asked to signal whenever they reached an insight. This approach would, on the one hand, allow the recording of the emotional reaction while interacting with the visualization and the data, and on the other hand it would offer a set of time intervals that each participant required to detect the above-mentioned pattern. After the task, each member of the first group was additionally asked to fill in a short questionnaire focused on their emotional reactions. Most users reported having no emotional reactions while interacting with the visualization. At the same time, four participants mentioned being surprised at some of the findings they made. However, the measurements from the EEG device suggested additional events with emotional aspects. Finally, six participants suggested that the EEG headset was uncomfortable to wear, while four expressed concerns about possible privacy issues (“You mean someone else can now see what I’m feeling?”).

Once the members of the first group finished the task and the questionnaire, the participants from the second group were given the same task, this time with the EmotionScents constructed based on previous user experiences displayed on the GUI of the visualization. Again, the time interval was recorded that each group member required in order to detect the pattern. Note that while the EmotionScents can offer information about the views or GUI widgets that were more often visited, this is only a secondary purpose of the representation. Compared to the work of Willett et al. [313], EmotionScents focus on offering information about the emotional states the users experienced during or right around the period of their interaction with a certain widget. If this information is not available or not relevant, the information about the frequency of visits is also not included (i.e., a white emotion bar does not allow the detection of gradients).

After analyzing the results from the evaluation, we found that the first group that did not have the aid from the EmotionScents took on average 310 seconds to find the pattern ($AVG_{noES} = 5\text{min } 10\text{sec}$, $SD_{noES} = 68\text{sec}$). At the same time, the second group managed to gather the same insights in 248 seconds ($AVG_{ES} = 4\text{min } 8\text{sec}$, $SD_{ES} = 44\text{sec}$). Furthermore, the statistical significance of the difference between the

two groups was confirmed in an independent samples t-test ($p = .015$). These values suggest that representing the emotions that users had during their interaction can offer new users orientational cues about application features that were emotionally remarkable to their predecessors.

To inspect the potential for self-awareness, the participants from the second group were also presented the simple Java compiler from Figure 4.32. After a scenario was presented to them that was similar to the one highlighted in the Use Case section, they were asked the utility of such a representation for self-awareness. In this case the results were more evenly distributed, as six participants could imagine having EmotionScents included in various interfaces for enabling emotional self-awareness, while the others were more skeptical.

Finally, the members of the second group were asked to fill in a questionnaire. This time, the questions were focused in the representation of EmotionScents and their utility. All participants considered that the color coding is useful, while five participants suggested that the size of the emotion bars is rather small not allowing them to see details. Still, after a informal inquiry, these participants agreed that they still had the possibility of detecting the different colors and their ratio in the bar.

4.4.5 Discussion

Our findings from the evaluation of EmotionScents suggest that this integrated, in-place visualization of user emotions can enhance window-based interfaces and offer additional information about the subjective impressions and experiences of users while interacting with the GUI. Also, the supplied emotional information incorporates two facets of the emotional space, namely valence and arousal, offering an overview of the emotional reactions of one or multiple users in previous or the current session. While inspecting the design guidelines, we can notice that the enumerated requirements have all been considered in the current representation.

In terms of integration, the GUI has suffered only minimal changes, and the location of the scents allows for good comparability. More importantly, the visualization encodes emotional information in a way that allows the detection of emotional trends on each UI component but without overloading the interface or the emotion bars. Additionally, the low resolution offered by a representation within a small display area

is also intended to reflect elements like: the complexity of human emotions, a limited precision of the acquisition of emotions, and a certain level of uncertainty in the correlation of user emotions and events (i.e., in some cases, the emotional states recorded at the moment when an event takes place might be generated by other environmental or user-internal aspects).

Note that the two visualization frameworks, EmotionPrints and EmotionScents, are very different in terms of representation and functionality mainly due to the different requirements and restrictions that multi-touch and GUI desktop applications have. In general, GUI applications revolve around standardized visual and functional constructs, called widgets, which should not be modified excessively in order to ensure recognizability and design consistency. At the same time, this restriction is less present in multi-touch applications, where the number of standardized widgets and applications that use them is far lower. Additionally, the difference in available dimensions for representing EmotionPrints and EmotionScents (e.g., color, shape, size) resulted in a different set of restrictions that the representations had to abide.

Furthermore, it seems important to highlight that one should not directly interpret emotion representations as indicatives of right and wrong. Both EmotionPrints and EmotionScents represent basically an additional level of information that is being collected and reflected, and that capture the user's subjective views. Having negative emotions in a certain case or when manipulating a certain UI component might in some cases even be desired, e.g., positive feelings after saving an image one has just modified in a photo editor, or negative feelings of frustration inside a strategic game.

Another vital aspect of recording and visualizing emotions in the interface is privacy. As one would expect, people are in general concerned about having their emotions recorded, distributed, analyzed. While some of these issues can be solved by collecting only anonymous data and, of course, with the consent of the user, there are also other techniques for ensuring user privacy. Currently, both visualization systems store the emotional readings in a database local to the machine on which the toolkit was installed. Therefore, the security and distribution of emotion data is entirely under the user's control. At the same time, in order to support data transfer and analysis,

users have the possibility of importing emotion information from external sources and represent them on a system where EmotionScents is running. As a general principle, however, emotional readings should be accessed, stored and distributed with the same privacy considerations with which any private information is managed nowadays.

4.5 Summary

In this chapter, we first explored the abilities of the Emotiv EPOC headset, a portable wireless EEG device, as an aid for recognizing user emotional states and employing them in the context of complex visual systems. For this purpose, we validated the device together with the proprietary framework in terms of detecting facial expressions and emotional states. We further explored the EPOC's functionality in two scenarios aimed at eliciting a set of emotional responses. Based on these tests, we highlighted out results as well as discussed possible pitfalls and disadvantages when using lightweight EEG technology.

Next, and more closely connected to visualization, in Section 4.2 we focused on moments of insight and their interconnection with emotional experiences. In a set of visual tasks, we compared the self-reported moments of insight to the affective readings offered by the EPOC device. The obtained results suggest a strong correlation between certain emotional states and Aha! moments, opening the door towards novel non-intrusive ways for detecting insight in visualization users and potentially offering alternatives in terms of evaluating visualization techniques.

After having highlighted our approach for the detection of user emotional states, the Sections 4.3 and 4.4 detail two emotion visualization techniques for representing affective states on multi-touch and desktop interfaces in real-time. First, for multi-touch systems, we have proposed a visual framework called EmotionPrints, that is aimed at increasing user emotional self-awareness and the awareness of other users' emotions in collaborative and competitive touch-enabled applications, both in real-time and offline. Besides the design guidelines and the implementation, we describe how the visualization of user emotions can support the evaluation of user experience and multi-touch applications. Second, we detail the rationale, implementation and functionality behind the EmotionScents, a framework aimed at representing real-time and post-task user emotions by enhancing pre-existing GUI widgets. Similarly to EmotionPrints,

this desktop implementation has the goal of increasing emotional self-awareness and the awareness of the collaborators' emotions in order to support the social interaction and the entire decision making process. Our two user studies suggest that both representations can capture and convey user emotions, while at the same time allowing collaborators to make more informed decisions and take control of their interpersonal interaction. A noteworthy fact is also that the proposed visualization frameworks are independent of the technology through which user emotions are recognized and reported.

After having presented our WebComets visualization system (Chapter 3) and our approach to detecting user emotional states and representing them for supporting emotional awareness and collaborative interaction, Chapter 5 of this thesis starts by addressing a second aspect of user-centered collaborative visualization, which is more focused on interaction in co-located collaborative scenarios around a tabletop display.

Chapter 5

Collaborative Interaction with Parallel Browsing Histories

In the previous two chapters of this thesis, we have addressed the topics of information visualization (representing multi-user multi-session parallel browsing histories with WebComets) and interpersonal interaction support (detecting and representing user emotional states with EmotionPrints and EmotionScents). In the context of this thesis, the WebComets visualization from Chapter 3 represents the backbone visualization that allows us to validate and evaluate our user-centered collaborative visualization techniques. While a set of these techniques has been presented in Chapter 4 and focused on emotional awareness and social interaction, this chapter addresses a second dimension of UCCV, namely fluid and flexible interaction on tabletops.

One of the main advantages of tabletop displays, besides their size and touch functionality, is their natural ability to support collaboration. Similarly to real world interaction when sitting at a table, tabletops allow users to gather around a horizontal touch-display, thus maximizing the number of operators in a way that not even large-scale displays can support. However, this is not the sole advantage of tabletops, especially in the context of co-located collaboration.

Over the last decades, tabletops have received an increasing amount of attention from developers, researchers and industrial users. The reason for this orbits around the interactive and social potential of these displays, a potential for which appropriate technologies and manipulation techniques are being explored. As such, one important attribute of the tabletops is their horizontal layout that allows various objects to be rested on their surface. In systems where these items can be detected and attributed a certain functionality, they are called *tangibles* or *tangible objects*.

Combining the tabletop experience with tangibles can enhance the experience of the users in multiple ways: tangibles can add a third dimension to the interactive representation on the display [17, 264], they can offer a seamless integration of concepts from the virtual and real world [177, 286], and—perhaps more importantly—they can be employed as physical metaphors that tap into the users’ real world experience, allowing them to minimize the learning curve of the system. In terms of metaphors, tangibles can be therefore divided into interaction metaphors (e.g., buttons [131, 287]) and visual metaphors (e.g., houses [108, 289, 320]), allowing the designers to draw on different areas of expertise. More importantly, however, these metaphors that are subtly encoded by the tangibles can also support the process of collaboration around a tabletop, by addressing issues like communication, feedback and adaptability.

This chapter is centered around a novel design for passive tangibles, entitled TangibleRings. TangibleRings are focused on overcoming some of the key drawbacks of many common tangibles by addressing issues like occlusion, tracking precision and in-object touch events. More importantly perhaps, these ring-shaped tangibles support flexible interaction and filtering in collaborative settings by enabling ring-in-ring nesting. After highlighting the design, implementation and functionality of the rings, we continue with a user study that both explores the advantages of using TangibleRings in a collaborative setting as well as evaluates the subjective dimension of the collaboration process. These studies are executed on a multi-touch adaptation of the WebComets visualization, entitled WebComets Touch. This second version of our visualization is a partial reimplementaion of the WebComets tool described in Chapter 3, offering broadly the same functionality as its predecessor, but with a range of supported interactions specifically designed for multi-touch and collaborative environments.

The following sections include elements from the publications entitled “TangibleRings: Nestable Circular Tangibles” [76] and “Visualizing Group Affective Tone in Collaborative Scenarios” [45]. Furthermore, parts of Section 5.2 are also detailed in the bachelor’s thesis of Felix Schmidt, entitled “Collaborative Motif Search in Large Browser History Visualizations through Combinable Tangible Rings” [247].

5.1 Nestable Tangibles for Collaborative Interaction on Tabletops

In recent years, the multi-touch functionality of tabletop computers has been augmented by the use of tangible objects that can be detected when positioned on or close to the surface of the tabletop. While these tangibles have the advantage of providing the user with distinct physical objects that can be used to interact with the virtual space of the tabletop, there are also some known issues. In order to be detected, tangibles are usually positioned on the surface of the tabletop, thus introducing a new set of challenges to the designers, e.g., display occlusion, reduction of touch-interactive space, limited flexibility in combining and interpreting the tangibles, etc.

In this section, we highlight TangibleRings, our low-cost solution for addressing some of these limitations that seem to affect many tangible systems (Figure 5.1). TangibleRings are a set of ring-shaped objects that reduce the display occlusion and allow users to inspect and interact with the virtual elements presented inside their perimeter on the tabletop. Furthermore, the rings can act—depending on the tabletop application—as the physical boundaries for virtual lenses (e.g., Magic Lenses [24], GeoLenses [297], FingerGlass [137]) or tangible views for information visualization [264], thus offering local views, filters and queries to the user. The tangibles also allow users to manipulate the filtered information by performing touch operations inside them, as well as store views or configurations to a particular ring. These features are supported by circular markers positioned on the bottom of each tangible, which have the ability to uniquely identify any *TangibleRing* placed on the tabletop.



Figure 5.1: TangibleRings in a map-based application, each ring controls different information layers. In this example, the single rings effect only their corresponding lenses, while manipulating the nested rings influences the view on the entire screen.

Furthermore, TangibleRings can be nested (i.e., concentrically combined) for increased flexibility in data manipulation and user collaboration, while also acting like overlapping knobs that allow the users to turn each ring separately. As such, TangibleRings provide users with the ability to execute logical operations on the rings in order to combine two or more queries or filters, while at the same time maintaining the advantages of minimal occlusion and in-ring touch operations. To support nesting, the rings have different diameters such that a smaller ring can be positioned inside the next larger one.

The contribution of the TangibleRings lies therefore in their ability to address multiple tangible-related issues at the same time:

1. *Occlusion*: Most block-like tangibles cover a significant area of the tabletop display, while at the same time offering a relatively low level of precision in selection and manipulation tasks.

2. *Reduction of touch-area*: Common tangible solutions do not only cover a part of the tabletop’s visible surface, but also block access to the underlying touch functionality.
3. *Dimensions of interaction*: Tangibles should enable new interactions by implementing visual or interaction metaphors, as well as supporting intuitive gestures.
4. *Collaboration and data manipulation*: Irrespective of the task, the tangibles should support the co-located collaboration process (e.g., aid in the distribution and management of private and shared tasks/views). Moreover, tangibles should support exploration, selection and manipulation of complex data, filters, controls, etc., by enabling functional aggregation.
5. *Cost*: Tangibles should be inexpensive in order to allow a wide range of applications and users to have access to them.

While other related solutions presented in Chapter 2 cover a subset of these concerns, none of them offers a similarly simple and flexible approach to addressing all these issues. In the following, we describe the design and functionality of our ring tangibles, as well as highlight their applicability in a set of examples. We then engage in a discussion about the advantages and limitations of our solution.

5.1.1 TangibleRings

TangibleRings are based on ring-shaped design [76] that offers a set of advantages compared to block-like tangibles. Firstly, employing rings with narrow boundaries and hollow interiors as tangibles reduces the level of occlusion introduced by these objects. Furthermore, rings can be positioned more precisely as users are able to inspect the virtual objects underneath them.

Besides occlusion, similarly to virtual lenses [24,137,297], the area inside the perimeter of each ring can be used to manipulate and display custom data through multi-touch interaction. While most tangibles hide part of the touch surface, through our approach users can still execute touch events on the tabletop screen regions situated inside the rings. These gestures are usually affecting only the information displayed inside the rings. For example, rings can be used as views acting like semantic lenses that display custom levels of detail for the information underneath them. Users can move these

rings in order to explore the entire space of the visualization, without interfering with the work of others. However, in certain applications, interior gestures can also be interpreted as global actions influencing the entire display. Furthermore, gestures executed immediately outside the rings' perimeters are mostly used to interact with the circular menus that are attached to each tangible [130].

One reason why tangibles are employed to augment tabletops is the lack of intuitive multi-touch metaphors for certain operations. For example, while there is a widely-employed touch gesture for zooming on a tabletop (i.e., pinch-and-zoom), a similar gesture in terms of intuition and acceptance for rotation does not yet exist. Consequently, many map-based applications on mobile devices do not even support rotation, whereas their desktop counterparts do. In such cases, through their third dimension and tactile feedback, tangibles can offer an additional playground for interaction designers when considering operations. Supported by their circular shape, TangibleRings can act like knobs and control buttons on the surface, thus implementing operations like orientation, selection and value adjustment. This interaction becomes particularly important when used on nested rings, where users can rotate each of the combined tangibles individually, thus manipulating different parameters for each associated filter.

A novel aspect of our approach is the ability to combine tangibles by nesting them. In terms of semantics, placing a tangible inside another creates a soft fusion between the two. With a similar logic like the motif search in Chapter 3, this compound tangible can now be used as a semantic combination of the functionalities previously attributed to the two individual rings. The complexity of the combination can range from a simple blend (e.g., display both visual filters) to more complex logical operations (e.g., filter out elements that comply with condition A associated with the first ring, but which do not comply with condition B associated with the second).

At the same time, through their lens-like shape, TangibleRings can support the distribution of private as well as shared tasks and spaces, enabling users to work both independently—for example, on a separate view represented inside or around one particular ring—and in the shared tabletop environment.

The final requirement to our design was reproducibility. Therefore, our aim was to avoid electronics and design a set of passive tangibles that individuals and research groups could quickly deploy as simple, low-cost solutions.

Design and Implementation

By choosing the form of a ring as the tangibles' shape for our approach, we had to start from scratch and prototype our own customized ring objects. For the early prototypes, we used slightly modified, ring-like real-world objects, that can be found in supermarkets, home-improvement markets, or toy stores. These objects revealed rough requirements for the needed diameters and minimum rim width for placing our markers. Based on these findings, we went on to the final prototyping step and modeled the rings in Autodesk 3d Studio Max¹. The resulting 3D data was plotted as a three dimensional solid object using a Dimension 1200 series 3D printer² (see Figure 5.2). Note that the height of our rings was designed to be inversely proportionate to their diameter to allow users to manipulate each ring individually when these are combined. Also, instead of modeling plain and smooth rings, we decided to create a non-slip surface in order to improve the usability of the tangibles.

After dealing with the reasoning behind the physical configuration of our tangibles, we had to address their detection on the tabletop surface. The tracking of our ring tangibles had to fulfill several requirements:

1. Detectable position and radius of each tangible,
2. Detectable orientation for each tangible,
3. Unique IDs for multiple tangibles,
4. Detection of concentric positioned tangibles, and
5. Interaction and gesture recognition inside the tangible.

For the implementation of the TangibleRings concept, we used the Microsoft Surface 2.0 platform. While the corresponding Microsoft framework already comes with marker recognition, these markers were rectangular in shape and unsuitable for our design. Thus, we had to think of a new marker design that works for our approach. As the Microsoft Surface gives access to the images of the built-in IR cameras, we decided

¹Autodesk, <http://www.autodesk.com> (February 2014).

²Dimension 1200es, <http://www.dimensionprinting.com> (February 2014).

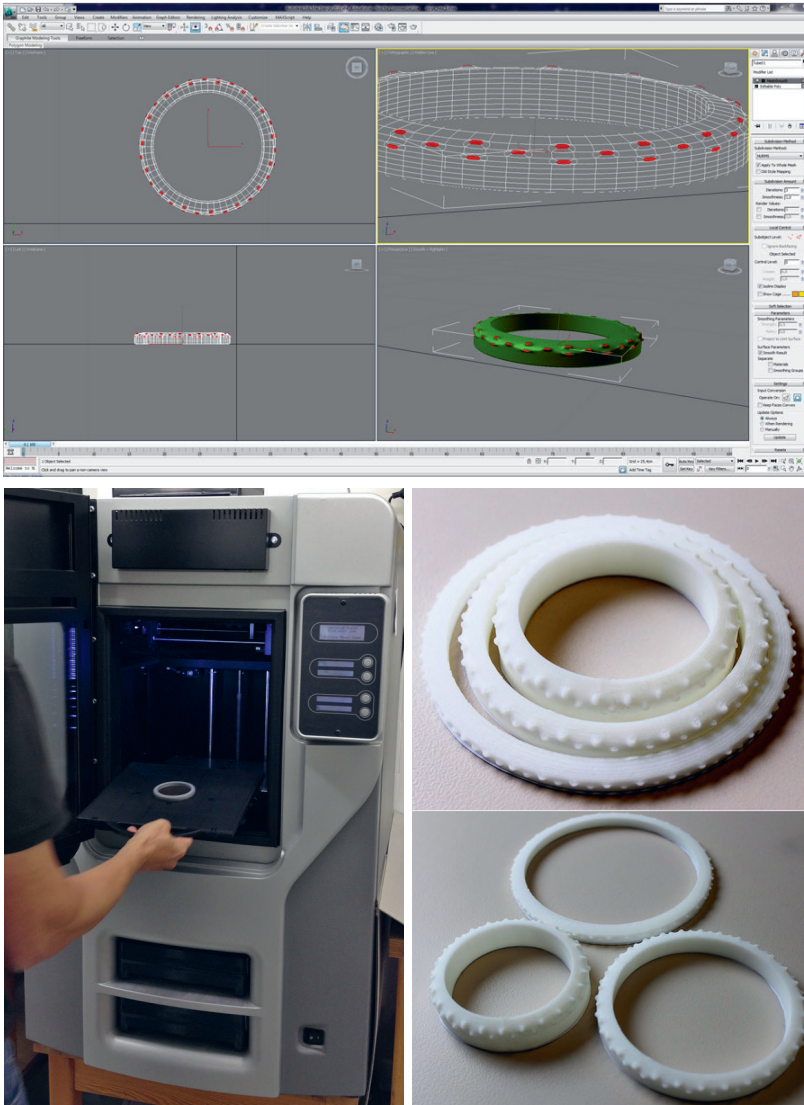


Figure 5.2: Prototyping the TangibleRings: 3D design of custom tangibles (top), prototyping the tangibles with a 3D printer (bottom-left), and final ring tangibles (bottom-right). Note that the height of the rings increases inverse proportionally with their diameter in order to allow users to manipulate each ring individually when they are nested.

to use OpenCV³ for the interpretation of raw images and the detection of our own custom markers. These markers had to fulfill several requirements: uniquely identify each tangible, support position and orientation detection, and allow the detection of nested (concentric) tangibles.




Figure 5.3: The detection of the correct pairs of ring markers. The green rings are the correct rings, the red rings are candidates that are eliminated due to having not enough matching id markers.

The TangibleRings markers are represented by a set of reflective blobs positioned on the bottom rim of the ring (see Figure 5.5). The idea behind these markers was to keep them as simple as possible to ensure detection stability. More specifically, to correctly detect each individual ring and its location (i.e., center and radius), a pair of elongated blobs was added to the rims of each tangible. The two circle detection blobs were positioned on opposite sides of the rim. The corresponding two points on the rings are enough to determine the center and the radius of the ring, while offering a quick and stable detection for one tangible.

However, when multiple rings are positioned on the tabletop, it might be difficult—even with the help of distance restrictions—to identify the correct pairs of circle detection blobs. In order to compensate, smaller blobs have been added to the rim at constant intervals, blobs that also encode the unique ID of each ring. These circular

³OpenCV (Open Source Computer Vision Library), <http://www.opencv.org> (February 2014).

blobs are about half the size of the detection blobs, and thus easy to separate by blob size. We used size and not color since this allows for a more robust detection. This also has the advantage that when thresholding the image for the blob detection, the upper threshold parameter is not needed anymore as the blobs are white.

In the case of multiple detected tangibles, the system considers every pair of large blobs as a candidate for a circle. Each potential circle is then validated by counting the number of small blobs that were assigned to the circle (Figure 5.3). If this number is larger than a threshold, then the pair is verified as a circle. This yields satisfying results for the following reason: during the detection, the small blobs are added to the best matching circle. In cases where only one or two small blobs have been assigned to a circle, it is still possible that these assignments have been made to a non-existent ring. However, each ring tangible needs to have at least five small blobs, thus avoiding false positives. Figure 5.3 shows this situation. The red markers can be assigned to a wrong ring candidate, but it is impossible to assign more than three ID markers to wrong circles. Mathematically, we need more than $2n - 2$ small blobs, where n is the number of ring tangibles, since every ring candidate can cut every ring two times minus the two cuts through the ring markers.

Once the correct pairs of circle detection blobs are established, we can easily compute the radius and the center (half of the distance between the blobs) of each `TangibleRing`. Additionally, since the size of a tangible does not change over time, we can use the knowledge about its radius and store it in a list together with its ID to improve stability. By a comparison of the detected radii of the rings and the radius associated with the ID, we can filter out false positive ring detection.

One problem when working with objects on the tabletop is that sometimes contact of the object with the surface is misinterpreted as finger touch event by the system. To overcome this problem we implemented a distance criteria δ (half the width of the used ring) and ignored touch events that lie in the area of $radius \pm \delta$ around the circle center in the area where the ring touches the surface. As such, the touching capabilities inside and directly outside the ring are maintained, e.g., for implementation of a menu around or inside the ring.

To further stabilize the detection, we implemented a state history that takes care of adding and removing rings from the system. If during one frame a ring is no longer detected—e.g., the user moves a ring and lifts it during this process—the ring is not removed immediately. Only if it is removed for several frames, it is deleted completely from the system, thus avoiding undesired artifacts.




Figure 5.4: Computation of ring orientation γ : to identify blob B as the marker for orientation, we compute the angles α_i between the markers on the ring and use the marker with the largest α . The orientation γ is finally given as the angle between \overrightarrow{CB} and the x-axis.

In terms of determining the orientation of a tangle, we employ the non-equidistantly positioned small blobs for each tangle. Let C be the center of the ring, as highlighted in Figure 5.4. We use one of the blobs B as the indicator and compute the angle γ of the vector \overrightarrow{CB} to the x-axis via the scalar product. However, for this purpose the system needs to uniquely identify a blob that can encode the orientation of the ring.

As the small blobs are positioned non-equidistantly on the rings, the orientation can be deduced by computing the line defined by the center of each ring and the blob B that has the largest gap to the next marker blob in counter-clockwise direction (see Figure 5.4). This keeps the coding simple and more robust.




Figure 5.5: Ring detection using custom markers for single and concentric rings: (left) The raw IR images obtained from the tabletop; (right) The corresponding images after filtering. These images are used as input for the blob detection algorithm. Large blobs encode ring position and radius, while small blobs encode ID and orientation.

To uniquely detect a ring, we detect circular shapes on the tabletop and reflective blobs representing ring markers. Afterwards, we combine these two findings by intersecting the detected rings and blobs. The recognized markers are then assigned to their corresponding rings. For this assignment we use the center of gravity of the blobs as the position of the marker. Then, we compute the distance of each marker to the center of each ring and assign it to the ring with the best matching radius, i.e., the difference between the radius of the ring and the distance of the blob to the center

of the ring has to be within a given range defined by a threshold. The ID of the ring is determined by the number of blobs on the ring. Additionally, in order to ensure accurate identification and filter out false positives, the system can check the pairing of the detected ID and the associated ring diameter in a dataset.

Due to the maximal reflective white color of the blobs, we are able to eliminate influences by fingers and hands of the users, shadows or bad lightning conditions by filtering of the raw image. This is illustrated in Figure 5.5, where the images on the left show the raw representation as obtained by the IR camera, while the ones on the right show the thresholded image for the blob detection. The lighter pixels from both the noise and the hands of the users can be filtered out completely, making the ID detection stable.

Another key aspect of our approach is the possibility to combine information related to the ring tangibles by nesting the tangibles themselves. Concentric rings are recognized independently, yielding two separate circles instead of a single large one. To establish if two or more rings are nested or not, we compare the distances of the found circle centers. If the distance between two circle centers is lower than the largest of the two associated radii, the circles are connected and each ring stores its partner(s) in a list sorted by radius. Additionally, such a list allows the system to know if a ring used in a nested configuration is an inner or outer ring, and it can act according to the application, e.g. only the outer ring should offer an ring menu on the outside like in Figure 5.6.

In terms of interaction, the ring-shaped tangibles can accommodate a variety of metaphors and gestures (see Figure 5.6). Besides the well-known translation and rotation of the tangibles, users have the ability to use these rings in both input and output scenarios. In terms of input, users do not lose the ability to interact with the touch surface inside the rings. Furthermore, the gestures executed inside the rings are detected accordingly and can be interpreted in a custom manner, e.g., manipulating only the information connected to the ring or exploring a subset of the system with the help of a physical lens. Additionally, the rings can be envisioned as control knobs that change the values of attached parameters through in-place rotation.

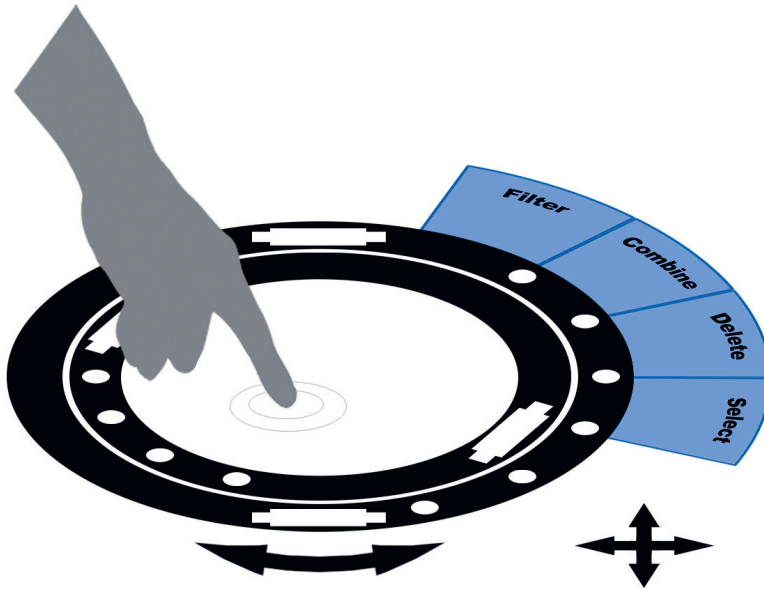


Figure 5.6: Interactions supported by the TangibleRings: translation, rotation (independent for multiple concentric rings), interior touch surface, and outside ring menu for application specific functions.

In order to control more advanced settings, a double tap inside the touch surface of a ring brings up an exterior menu, similar to the facet wheel in [130]. Through this menu, users have the ability to select parameters that they can manipulate by rotating the ring clockwise or counterclockwise, or by gestures inside the ring area. Also, the menu enables users to store their filters, views, queries (depending on the application) to the ring. More precisely, the parameters of the query together with the ID of the ring are stored into a database on the tabletop.

After a ring is removed, that particular configuration can only be obtained by repositioning the ring on the tabletop screen or by recreating it from scratch. This allows users to exchange rings or remove rings from the surface until the next time they want to use that view. Interpreting removed rings in such a way also addresses some permission and privacy concerns, by allowing the rings to emulate the functionality

of a key that unlocks preset configurations and stored data. In terms of output, the rings support the lens metaphor that allows the tabletop system to display customized data inside or outside their radius. These and other interaction capabilities are further highlighted in Section 5.1.2.

5.1.2 Application

For demonstration purposes, we have developed an interactive map tabletop application for exploring urban areas. The application uses our TangibleRings concept to manipulate a set of features (see Figure 5.7). At the core is a satellite map of an area that can be augmented by the users with new layers of information. Each ring can be assigned one or a combination of these layers that show different information inside or outside the corresponding tangible (e.g., map view, traffic or population density, air pollution, street names, heat map, etc).




Figure 5.7: Map-based application that supports TangibleRings and runs on a tabletop.

Figure 5.1.2 shows three rings representing information about street names and traffic density, as well as a heat map. By combining the corresponding rings, a user can visualize multiple layers at the same time, usually restricted to the interior of the nested rings. This is shown in 5.1.2. The rings are combined and so is the information

within, giving the users the chance to gain new insights by exploring the resulting view. Note that while nesting two rings results in the combination of the two associated layers, there are applications where more complex inter-tangible logical operations can be executed.




Figure 5.8: Single rings showing layers with street names, traffic density, and the urban layout (left). When combined, the two rings show the traffic density and the urban layout inside their perimeter (right).

Additionally, when employing rings as virtual lenses, users have the option of zooming and locking a view (see Figure 5.9). By using in-ring gestures or rotating the rings when the zoom parameter is selected, users can explore the dataset locally without affecting their collaborators. By default, rotating the tangibles inside the map application results in a manipulation of the opacity parameter for the corresponding layer. However, the parameter that is controlled by the rotation of the rings can be modified, thus also supporting operations like zooming, changing color schemes, and distance measurements.

The rings allow users to select two lock functions by means of their ring menu (see Figure 5.9). One of them focuses on view rotation, and is meant to lock the orientation of the view inside the lens to the orientation of the ring. When this setting is on and a user rotates the ring, the view inside the tangible follows the rotation of the ring. This allows users to readjust the orientation of their lenses in order to read text more easily or compare information with different angles. The second lock function involves view translation. Users that find relevant information in an area of the map can decide to lock the lens to that position in order to display it in another context or simply share it with other users by passing them the tangible. At the same time, a translucent halo



Figure 5.9: An example of TangibleRings enabled interaction: (top-left) the rotation of the ring controls the opacity of the information inside the ring; (top-right) locked view displaced and given to another user to inspect; (bottom) the zoom factor can be changed by using the ring menu or by rotating the ring after selecting 'zoom' as the new control parameter.

is being displayed at the original location where the view has been locked in order to allow users to make the correspondence. Moreover, users have the ability to execute touch gestures inside the ring with the locked content in order to further explore the surroundings of the view. These operations result in an update of the location of the view-origin halo on the tabletop.

Furthermore, when considering collaboration and level of clearance, TangibleRings have the distinct advantage that users can remove their rings from the tabletop in order to make their custom lenses or filters inaccessible to others. Custom information is thus stored to the tangibles, which can be taken from the tabletop and cannot be accessed again without the physical token. When the ring is repositioned on the tabletop surface, the information layers will be restored for further analysis.

In cases when users wish to share their findings with the entire group, they have the option of displaying their custom layers inside, outside or on the entire screen of the tabletop. This is achieved through the ring menu of the corresponding tangible. Irrespective of where the information associated with a tangible is displayed, the user still has the ability to manipulate that layer through ring rotation, touch actions inside the perimeter of the ring or through the ring menu. All these functionalities also support a seamless transition between collaborative and individual work, such that users can maintain an optimal level of flow, as suggested by the characteristics of fluidity highlighted in [81].

5.1.3 Discussion

While in the previous segments of Section 5.1 we highlighted the advantages of our approach, there are also a couple of issues raised by TangibleRings. In terms of the physical design, it is obvious that due to the different diameters of the nestable rings, we can only deploy a limited number of tangibles at a time. Although we plan to increase this number by reducing the width of the rings, it is clear that this limitation will persist. However, we argue that due to the capability of combining rings and transferring filters or constraints from one tangible to another, users will always be able to free up tangibles for novel operations. Thus, the number of rings on a tabletop seems to be more strongly connected to the number of collaborators, as having two rings per user already allows for the composition of very complex views and constraints.

Besides the limitation to the number of rings due to their varying radius, users have to also consider the nesting restrictions. While two non-consecutive rings, in terms of diameter, can be combined and detected, the use of such a configuration is quite application-specific. Although some users in the study tried to combine a larger and a much smaller ring in order to use the small one as a joystick inside the area defined by the large one, we do not suggest this combination in systems that require precision positioning or knob-like rotation of nested rings.

Furthermore, our current implementation does not consider possible touch events between the sides of non-consecutive nested rings. With the presented marker configuration, problems can also arise if the color of a ring in the IR image is similar to a hand in its proximity, resulting in a large blob instead of a circular shape. This is, however, usually avoided through filtering and color thresholding.

5.2 Supporting User-Centered Collaborative Visualization

A multi-touch implementation of the WebComets visualization described in Chapter 3 has been developed for co-located collaboration around tabletop displays. This visualization, a redesign of the one presented in [50], is aimed at supporting collaborative inspection of parallel browsing behavior on touch-enabled large displays. In Section 5.2.1, this novel approach to the WebComets visualization is described, focusing on a couple of changes in terms of representation. Additionally, we highlight a set of interaction metaphors that integrates the functionality of the TangibleRings in order to support flexible collaboration and seamless switching between private and shared interaction. The resulting system is called WebComets Touch, and it incorporates an adaptation of the EmotionPrints visualization from Section 4.3 and interaction techniques based on the TangibleRings from Section 5.1. The purpose of this visualization system is to allow an application-oriented evaluation of the user-centered collaborative visualization techniques proposed in this thesis.

5.2.1 WebComets Touch

The visual interface of the WebComets Touch is a reimplementaion of WebComets that is specifically modified to support touch-based interaction. The reimplementaion was developed on a SUR40 tabletop from Samsung⁴ and employs the PixelSense technology to aids the recognition of tagged objects. The application was implemented

⁴Samsung 40 SUR40 SMART Signage, <http://www.samsung.com/uk/business/business-products/smart-signage/specialised-display/LH40SFWTGC/EN> (February 2014).

in C# and WPF⁵ .NET, as imposed by the Software Development Kit (SDK)⁶ for PixelSense.

Tabletop users now have the option of opening multiple parallel browsing session, either in the same view or in different views. This is supported by independent session windows (Figure 5.10) that incorporate the visualization that would have been previously displayed as a desktop application window (Figure 3.2). Each user can open multiple session windows on the tabletop and load multiple parallel browsing logs in them. Additionally, the session windows allow users to pan, rotate and zoom their content, thus enabling users to efficiently share the tabletop display. If users want to analyze a dataset privately, this can be achieved inside one or multiple session windows that each of them manages. However, if at some point it is desired to collaborate on a representation, the same users can resize the windows and position them in such a way that cooperation is enabled.




Figure 5.10: The WebComets Touch visualization running on a tabletop. Interaction is supported through touch gestures and TangibleRings.

⁵Windows Presentation Foundation, [http://msdn.microsoft.com/en-us/library/ms754130\(v=vs.110\).aspx](http://msdn.microsoft.com/en-us/library/ms754130(v=vs.110).aspx) (February 2014).

⁶Microsoft Surface 2.0 SDK, <http://go.microsoft.com/fwlink/?LinkId=201003> (February 2014).

One important visual change in WebComets Touch is the representation of the tabs. While still horizontally encoded, tabs are now represented as rectangular tracks that run parallel to the timeline, from the moment a tab was opened and until it was closed (see Figure 5.11). This encoding is similar to the visualization of multiple heterogeneous simulation runs in World Lines [305] or to the temporal structure displayed by the History View in [268].

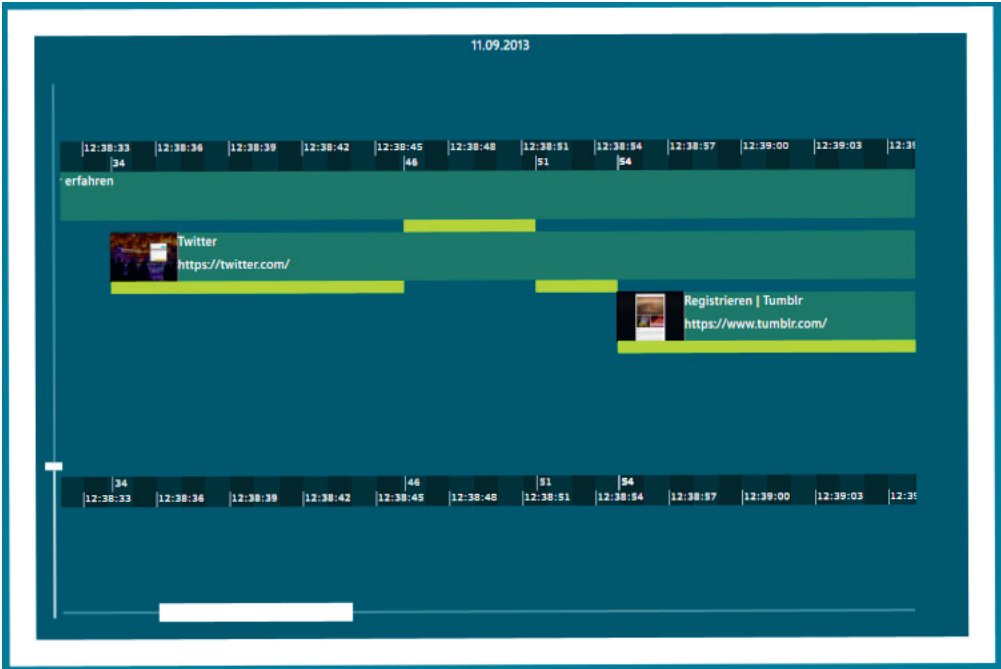


Figure 5.11: Window presenting a user’s browsing session. Each dark green rectangle represents an open browser tab, while the light green bars highlight which tab has been in the foreground.

Each tab is now subdivided into rectangular sections that represent the time interval when a particular website was visited. If a tab was open but no website was loaded, the rectangle would have a different color until the moment when a web page was opened. Moreover, users can now opt to display the various time intervals (focus and active time) as smaller rectangular segments underneath each website, also encoding the exact time intervals when these web pages were in the foreground. If a focus time bar disappears or reappears before a website was closed on a particular tab, it

means that the user selected another tab, window or application that got moved to the foreground, thus pushing back the previous tab (see Figure 5.11). Finally, due to the increased available screen area, the favicons have been replaced by screenshots of the web page that has been accessed, in order to improve recognition rates.

The redesign of the tabs, websites and the corresponding time intervals was motivated by the need to obtain a representation that is adapted to the requirements and specifics of multi-touch interaction (e.g., selection of narrow objects on touch interfaces can be problematic). As such, the implemented changes increased the size of various elements and separated their location in space, allowing users to more easily select and manipulate the various components.

In terms of touch interaction, users can execute touch gestures for all the selection and manipulation operations that were mouse-driven in the desktop visualization. Additionally, the interaction with the visualization is supported by the TangibleRings (see Figure 5.10). The function of the rings is mainly related to the motif search, where rings aid the creation of motifs/constraints, store them, enable their combination through the ring nesting metaphor, and control to which session windows the queries are applied.

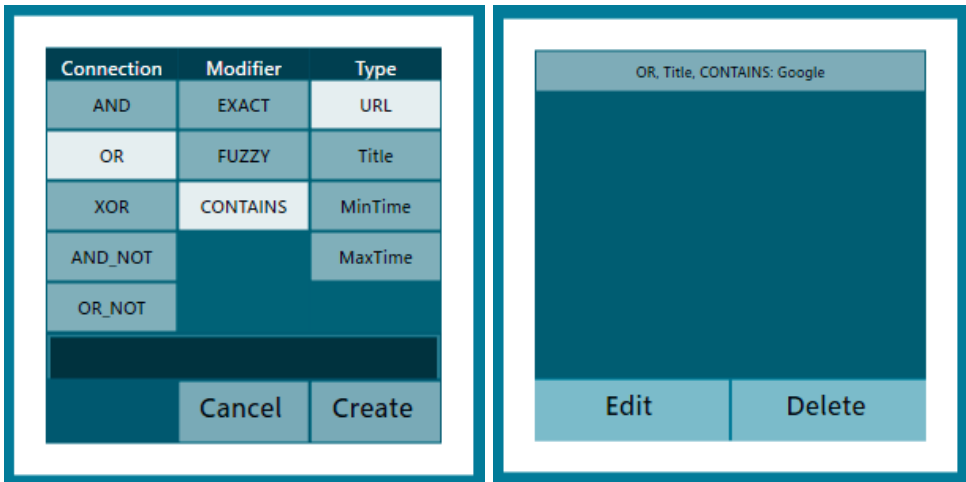


Figure 5.12: The constraints for a ring can be managed by accessing the following windows through the ring menu: constraint generation associated to a ring (left), constraint list for managing the constraints associated to a ring (right).

Initially, each ring can be attributed a set of constraints that represent the root elements of a motif (see Figure 5.12). These constraints can refer both to structural information (i.e., online navigation patterns like visiting six websites in a tab) and content data (i.e., which website was actually visited) related to the accessed websites, similarly to the description in Chapter 3. The combination of the constraints into more complex motifs is supported through logical operations, like AND, OR, XOR, and NOT (Figure 5.12). Additionally, similarly to the initial WebComets motif search, users can also compare time intervals and search for substrings in the browsing logs.

After creating all the required constraints, the corresponding motif is attached to the ring and its ID until it is moved or deleted. Furthermore, users can remove a ring from the tabletop, making the corresponding motif inaccessible to the other users. In such cases, the rings can be seen as access keys to the custom motifs that each user has created or modified. Moreover, users can manipulate their motifs through a ring menu similar to the one from [130]. This ring menu (see Figure 5.10) allows users to modify their motifs, apply them to one or multiple session windows, or to move and combine them through nesting. A motif search is applied to a session window by moving the relevant ring over the representation and opening the ring menu. The elements of the menu are dynamic, as the menu incorporates additional specific commands when the ring is located over a window.

A motif search that is applied to a session window is represented by a thin line that connects the ring with the representation. This offers in a high level of flexibility, as one ring-based motif search can be executed on multiple parallel browsing logs, while at the same time multiple rings can query one session window. The search results are visualized by highlighting the corresponding nodes, i.e., which fit the motifs, in the session windows. In cases where multiple queries are executed individually on the same representation, the results are highlighted with different colors, where each color corresponds to the one of the connection lines between the window and the rings. Furthermore, when multiple users decide to explore a session window collaboratively, they can enlarge the representation and employ their rings as lenses that only highlight nodes that fit the motif if these are inside their perimeter.




Figure 5.13: Two nested TangibleRings in WebComets Touch visualization. Each ring has a set of constraints attached to it, and the combination of these constraints is being applied to the two session windows to which the rings are connected by lines.

Besides storing search motifs, rings can also be nested in order to logically combine these motifs (see Figure 5.13). These combinations take place based on the same logical operations that have been previously enumerated. Moreover, when two rings are nested, users have the option of merging or moving a motif from one ring to another, thus resetting the initial ring. At the same time, the query connections to session windows for the initial ring would also migrate to the second ring. As already highlighted in Section 5.1, reusing TangibleRings by moving associated information is a vital process, as it enables users to experience flexible interaction with a limited number of TangibleRings.

5.2.2 Group Affective Tone Visualization

Until this point of the current chapter we highlighted the configuration of our tabletop system in terms of the WebComets visualization and the corresponding interaction metaphors that are being supported both through multi-touch and our ring-shaped tangibles. One feature that is still missing from our application is the incorporation of the subjectivity dimension. In other words, our desire was to incorporate a feedback module for emotional awareness in collaborative scenarios based on the techniques covered in Chapter 4, as well as evaluate the effect and usefulness of such a module.

While the EmotionPrints visualization from Section 4.3 might seem ideal for this purpose, research in the area of affect [93] as well as our preliminary tests suggested that employing a visualization that would clearly identify the owner of each affective information might be counterproductive, as the collaboration process is likely to be brought to a halt when users start to focus on individual contributions. However, one element that can suggest a healthy collaboration process is the *group affective tone* or GAT [94,95]. The group affective tone is defined as the consistent emotional reactions of the members of a group. While not all groups have an affective tone, positive GAT has been linked to higher effectiveness and creativity during collaboration sessions [63, 123]. Moreover, users “in a happy state is most likely to accomplish their pre-set goals, and will perform better than if they were in a negative emotional state such as anger or depression” [166].

To leverage on these findings, we introduce a novel visualization technique for representing group affective tone in collaborative applications and visualizations. The goal of this representation is to improve user awareness of the overall affective tone with the purpose of supporting the collaboration process by allowing users to react to divergent or negative affective tone levels.

In order to represent the group affective tone, we first need to interpret the users’ emotional states. As described in Chapter 4, the EPOC EEG headset together with Russell’s model of affect [230] was employed in order to generate the valence-arousal space of each user. From this data, only the user’s affective valence (i.e., positive emotions or negative emotions) was considered for the computation of the GAT.

In general, affect visualization approaches include abstract representations [184], avatars and affective icons [93, 121, 164], as well as custom interface widgets [164]. While some of these techniques are also focused on emotional self-awareness, there are few that consider the particularities of representing emotions in collaborative visualizations. In the case of representing the GAT on a tabletop display, the visualization would need to encode two levels of information: presence of GAT and, if present, the type (positive or negative GAT). Moreover, such a representation would have to consider the specifics of collaborating on a tabletop, where objects need to be accessible and interpretable from all positions around the table.

A mapping between the one-dimensional valence space, normalized to $[-1, 1]$, and a segment of predefined frequencies was generated. As such, each valence value in the domain $[-1, 1]$ will be uniquely connected to a certain frequency. These frequencies are then used for generating sine waves, such that for each valence value we define a frequency and thus also a unique sine wave. The sine waves obtained from all the users that are participating in the collaborative session are then added up, generating a new signal. Based on the summed sinusoids and the superposition principle, the resulting signal may or may not have a shape that is similar to a new sine wave. If the frequencies of the individual sine waves are similar, the resulting signal will be close to a sinusoid. Otherwise, the signal shape will be less homogeneous (see Figure 5.14).

To obtain an abstract representation of the GAT that is accessible to users situated at any position around the table, the final signal that results from the sum of the sinusoids is mapped to the contour of a circle (Figure 5.14). In this representation, the homogeneous contour suggests the presence of a group affective tone. Furthermore, the existence of such a tone is also captured by the saturation levels of the representation, as desaturated colors suggest the lack of an emotional consensus. At the same time, the colors red and green are present in cases when a group affective tone exists and they encode a negative or positive valence, respectively.

Besides improving the awareness of the group affective tone, the GAT visualization allows team members and leaders to increase awareness of the group's mood, to better evaluate the flow of the collaboration and to pro-actively engage in conflict management. As such, the GAT visualization approach can be employed in a wide set of collaborative applications and visualizations, especially in scenarios where unanimous




Figure 5.14: Sum of multiple sine functions (top) and the GAT representation for their mapping to a circle contour (bottom). The examples represent three distinct situations: (left) sum of identical low frequency sinusoids encoding a negative group affective tone; (center) sum of sinusoids with varying frequencies suggesting the absence of a group affective tone; (right) sum of identical high frequency sinusoids encoding a positive group affective tone.

decisions are vital. At the same time, one needs to consider that achieving a positive GAT does not ensure successful collaboration, as the excessive focus on obtaining a positive group tone can even distract team members from the task they are performing.

Figures 5.15 and 5.16 highlight the configuration of the final system. The WebComets visualization is being executed on the tabletop, while aided by the interaction metaphors offered by the TangibleRings. The collaborative system is augmented by the presence of the GAT visualization that offers affective feedback to the group in the form of convoluted emotional valence. The emotional readings are obtained through the information supplied by the EEG headset that each user is wearing.



Figure 5.15: Users working collaboratively on the WebComets tabletop application while wearing EEG headsets. The visualization incorporates the GAT representation.

5.2.3 Evaluating the Impact of User-Centered Collaborative Techniques

After highlighting the various interactive and visual modules of our collaborative visualization system, in this section we will present a set of studies that address these modules and their efficiency in collaborative contexts. More specifically, the next user studies focus on evaluating following components:

- TangibleRings as a technique for supporting collaborative interaction on tabletop displays,




Figure 5.16: Representation of the group affective tone cue during the interaction with the tabletop display.

- emotion visualization, specifically group affective tone visualization, in the context of increasing users' emotional awareness of subjective aspects, and
- an evaluation of a collaborative visualization through subjective measurements.

User Study: TangibleRings

A study was devised in order to capture the potential advantages as well as drawbacks of the TangibleRings approach. The study involved 12 participants (9 men, 3 women) with an average age of 22.6 years. They all had a background in Computer Science, but only five of them had any experience with tabletops and none of them had worked with tangibles before.

The scenario was developed around the collaborative browser history visualization described in Section 5.2.1. The participants were given the task of finding all the matches for two different browsing patterns, as well as the incidence of their combination. Each such three-task assignment would be executed on data obtained from multiple browsing sessions of up to four users. The participants were informed about having the option of exchanging information and collaborating on their task.

After a 60-minute introduction to the browser history visualization, the tabletop functionalities, and the features of the rings, the participants were divided into four groups of three. Each group was asked to perform two such assignments by using the system once with and once without the TangibleRings, in random order. In instances where the groups had to solve the tasks by using the browser history system only by means of the touch interface the access to the same features and functionality was ensured (e.g., session windows, constraints management windows, etc.). Also, for the instances where the participants did not use tangibles, constraints were grouped and represented as rings on the tabletop. Similarly to the physical rings, inside these virtual lenses/rings the users had the option of inspecting the currently active constraints and their logical combination inside these virtual lenses.

Each group was able to detect all instances of the occurring patterns. Moreover, in the assignments where the groups employed the TangibleRings the completion times were faster ($AVG_{noTR} = 29\text{min } 21\text{sec}$, $SD_{noTR} = 3\text{min } 42\text{sec}$, $AVG_{TR} = 21\text{min } 52\text{sec}$, $SD_{TR} = 3\text{min } 2\text{sec}$). Additionally, a paired samples t-test was executed, confirming a statistically significant difference ($p = .02$). The assignments where the physical rings were used presented an average speed-up of 25% compared to the software counterpart. The difference was especially visible in the third task, where the detection of the combined patterns with the help of TangibleRings would be executed 37% faster on average.

To further put these findings into context, we observed that when working with TangibleRings, group members exchanged their rings multiple times, while also explaining what they had found by positioning their rings over the session window of the other user. However, this was not the case with the software solution, although this func-

tionality has been previously highlighted to the participants. Note that all the groups have employed the rings both for private and shared exploration and interaction, at one or another point in time, with users often switching between independent (local view) and collaborative work.

The participants also were given a post-task questionnaire inquiring about the design of the tangibles, their utility and the corresponding interaction metaphors. The lens and nesting concepts were very well accepted, while the knob-like rotation raised a couple of question marks. One user was unsure he could rotate the rings in the same position without translating them. Another user said that he could rotate each individual nested ring accurately, but that he encountered problems when trying to separate the nested formation. In such cases, he was unable to grab any middle or even interior ring to lift it up without disturbing the rest of the formation. All participants stated that they could imagine using *TangibleRings* for other collaborative tabletop scenarios. As such, our study suggests that users comprehend the metaphors behind the ring interaction and are able to execute complex manipulations with relative ease after a training period.

User Study: Group Affective Tone Visualization

In order to explore the effects of representing the group affective tone during a tabletop collaborative session, a second study was executed with another set of 12 participants (7 men, 5 women) divided into four groups (see Figure 5.15). All the participants had a background in Computer Science, with nine of them having previously worked with tabletop displays. The average age of a participant was 21.3 years.

Similarly to the previously highlighted study, a scenario was designed that was centered around the *WebComets* visualization for tabletops described in Section 5.2.1. The groups had to search for two specific configurations of browsing patterns, as well as a particular combination of the two patterns (i.e., pattern A but not B), resulting in three search tasks grouped as an assignment. Each group had to execute two such assignments, once with the GAT representation and once without it, in random order.

Before commencing, the groups were given an introduction to WebComets and TangibleRings. Additionally, the specifics of the GAT visualization were highlighted, allowing them to extract and correctly interpret this additional level of affective information. Furthermore, the tabletop was augmented by an additional button that participants could press in order to report moments when they perceived the displayed GAT representation as conveying incorrect information. All groups were informed about having the option of communicating and collaborating on their assignments, however without suggesting that collaboration would be the focus of our study.

All the groups were able to successfully complete their assignments, although one group did require assistance at one point with the composition of the patterns. While no significant differences were noticeable in the average completion times of the assignments, the groups agreed that the GAT visualization has raised their awareness of the group's emotional state. For tasks where the GAT visualization was available, users communicated more and actively engaged in supporting the collaboration process (e.g., "Does everyone agree to proceeding this way?"). One explanation for the lack of an improvement in terms of completion time is given by the relatively low complexity of the assignments—which, while challenging, could also be addressed individually—and the time that has been added to each GAT-assisted group through communication about the team state and dynamics. Moreover, groups that were given feedback through the GAT representation were more likely to produce a team leader.

In terms of the accuracy of the GAT representation, a post-task questionnaire and the instances when users reported false GAT readings (e.g., positive GAT when a user experienced a negative emotional state) were analyzed. The results suggested that the GAT visualization was perceived most of the time as accurate, with only three instances where users pressed the report button in the entire study. When addressing these moments directly, users felt that their emotional state did not fit with the one transmitted by the representation ("I was already getting frustrated at that point.").

In a post-task questionnaire, participants were also directly asked about the GAT representation and its utility. Most users have found the visualization to be intuitive, allowing them to quickly discern the presence of a positive or negative group affective tone. One issue that was raised was related to the positioning of the representation on the display, as three users felt that they might not notice any changes if the circle was at the other margin of the tabletop. In terms of utility, ten participants could

see the benefits of using such a GAT visualization for increasing emotional awareness and supporting the collaboration process in large display scenarios. A positive impact of the emotional awareness was also noticed in the analysis of our video logs and the BCI-based subjectivity data, as described in Section 5.2.3.

Evaluating Collaborative Systems through Subjective Measures

Evaluating visualization systems [41], and especially collaborative ones [65], can be a very difficult task. At the same time, however, subjectivity readings have the potential of augmenting the information obtained through quantitative evaluations by exploring cognitive and social dimensions of the collaboration and interaction. More specifically, EEG measurements have been used in recent years to investigate user subjective aspects, like cognitive workload [5,6,20] and emotional states [211,212], and to employ the gathered information as an additional level of insight in evaluation scenarios.

In the context of the study from Section 5.2.3, with the same participants and tasks, we inspected the subjectivity data in order to draw conclusions about the visualization, interaction and collaboration aspects. We examined the changes in the group affective tone, but also the variations in the valence and arousal levels of each individual, and correlated these readings with the focus and actions executed by them, as captured through video recordings. Note that for our evaluation, although available, we did not consider the direct correlation of interaction location and emotion state of the users, e.g., EmotionPrints.

Figure 5.17 shows the GAT of a group and their corresponding individual valence levels during an entire assignment. By employing such representations, we were able to pinpoint a set of relevant moments in the affective data (e.g., abrupt changes). After inspecting the video logs for the corresponding time segments, we could sometimes establish a theory of what the action-reaction dyad might be in that instance. For example, in case of the group that is highlighted in Figure 5.17, we noticed how frustration—or more generally negative emotions—set in during the third task of the assignment and resulted in a negative GAT. These findings were confirmed by the members of the group that highlighted in their questionnaire that during the third task they were frustrated. Furthermore, in this instance, while raising awareness, the GAT visualization did not persuade the users to take action and improve their




Figure 5.17: Representation of the valence fluctuation for the three participants from one of the four groups in this study. These individual valence levels were not available to the participants during the tasks. The thick bar at the bottom encodes the GAT throughout the task, with red representing negative GAT, green being positive GAT and grey the absence of a GAT.

emotional states. The positive valence from the final minutes of the assignment has been mostly attributed to a moment of insight that one participant experienced when she discovered a novel way to combine the two patterns (see Figure 5.17, continuous line).

Similarly to the group and assignment considered in Figure 5.17, we analyzed the affective data generated by all groups in all assignments. Based on this information, we generated a set of potential correlations between user activity and the values of the subjectivity readings. We then cross-referenced our findings with the information extracted from individual informal discussions and a post-task questionnaire. The following presents a summary of our findings that have been supported by participant self-reports, the questionnaire and/or the video logs:

- Throughout our study, the nesting of TangibleRings has resulted in six distinct instances where fast motions or incorrect use have decreased user emotional valence, resulting in levels of frustration. More specifically, users that would move the rings fast over the tabletop’s surface would get frustrated over a lagging representation. Further, some users tried to nest rings of non-consecutive sizes, a functionality currently not supported by our system that resulted in a lack of responsiveness (e.g, ring menu).
- Three users also experienced negative emotional states most likely due to the frustration induced by the manipulation of the WebComets windows. The windows of the tabletop version of WebComets present a rim that allows users to drag and resize them. Still, under some lighting circumstances—including the effect of users moving around the tabletop—the touch precision on these narrower rims can decrease. This represents useful information for a novel iteration of the WebComets implementation, information that was only visible in one of the three corresponding video logs. Moreover, this information might have been omitted in the self-reports if it would not have been addressed directly by the interviewer.
- In terms of collaboration, we could notice a couple of instances where interpersonal dynamics have influenced (i.e., changed) the emotional state of the group members. The influences took the form of both negative and positive changes in the dimension of valence and arousal. For example, in one instance a participant was contradicted by the rest of the group, resulting in a drop of emotional valence and increase in arousal. At the other end of the spectrum, intense discussions about the tasks and the features of the system have resulted in instances where the valence of all the involved users increased.

Additionally, we were able to detect two instances where participants experienced moments of insight. By inspecting both the levels of frustration and excitement, we identified these instances where the emotional activation patterns would fit the one described by the method in Section 4.2. By cross-referencing the information with the images captured by the video logs, we could conclude that both users had a moment of deeper understanding. While one of the Aha! moments was related to the combined pattern that needed to be composed for the third task, the other one was

actually related to the interaction with the TangibleRings (further supported by the user’s exclamation “Oh, you can move it like this?!”). Moreover, during an informal debriefing and while being presented with the corresponding video sequences, both participants confirmed experiencing something that could be called an insight.

Note that however precise our interpretations of the video feed and the affective information are, the correlations between the user perception and physical actions and the emotional reactions are speculative in nature. As widely highlighted in literature [241], emotional states can be interpreted as reactions to both exterior/evoked (e.g., system, interface, social and physical environment) and interior/induced (e.g., thoughts) actuators (see Chapter 2). Under these circumstances, the correlations highlighted above—as well as the ones of any real-world action–user emotional reaction dyad—should be perceived only as potential cause–effect couples. However, we argue for the fact that given a large enough number of participants, subjectivity readings obtained through non-invasive and mobile technologies can offer feedback and probable correlations of how the users’ emotional states evolve together with their perceptions of the system—be it visualization applications, their interaction or the collaborative context.

5.3 Summary

In this chapter, we first presented a new kind of passive tangibles, entitled TangibleRings. Our design employs ring-like objects of various diameters in order to support interaction with the tabletop area situated inside the rings, reducing the level of occlusion and allowing rings to be nested. Our approach can be easily reproduced and is affordable, as TangibleRings are based on a simple design combined with custom circular markers. Furthermore, the nested tangibles offer an intuitive way to create and combine constraints to generate filters, views and queries, as well as support both private and shared work on the tabletop.

The features and functionality of the TangibleRings are highlighted in detail on an adaptation of the WebComets visualization that runs on tabletop displays and is designed for co-located collaborative exploration and analysis. In this context, we draw on our work in detecting and visualizing user emotions (Chapter 4), and propose a novel approach towards emotion visualization and awareness for collaborative scenarios by computing and representing the group affective tone (GAT) of the group working around the tabletop.

Finally, a set of user studies offer a unified view of the research techniques covered in Chapters 4–5 and highlight their importance for user-centered collaborative visualization. These user studies focus on: the TangibleRings as a technique for supporting collaborative interaction on a tabletop, the GAT visualization for increasing users' emotional awareness in collaborative settings, and an evaluation of the entire system through subjective measurements. Our results suggest that TangibleRings and the GAT visualization support the social and interactive dimensions of tabletop visualizations by considering the specific characteristics and needs of both the individual users and the group as a whole.

Chapter 6

Conclusions and Outlook

The following sections offer an overview of the research included in this thesis. First, a critical discussion addresses the research described in Chapters 3–5, focusing on the various advantages and pitfalls of our proposed solutions. Next, we expand on the theory behind UCCV with a set of properties formulated based on our research experience, and highlight the social implications of our work. In Section 6.2, we summarize our work and highlight our contribution. We conclude by presenting potential future research related to our approaches.

6.1 Discussion

Focusing first on parallel browsing log visualization, in Chapter 3 we introduced WebComets as a solution that allows analysts to inspect, explore and search in large online navigation logs. As mentioned in this context, one issue that our approach had to tackle was scalability. While it has been addressed through the various interaction techniques and views (e.g., zooming, branch collapsing, highlighting, or motif searches), scalability still remains a concern. Especially in cases where highly parallel browsing histories are loaded, it is difficult to explore the entire width of the representation. Even if collapsed branches help in this sense, future improvements of the visualization will consider a more flexible and user-defined approach to branch positioning, such that analysts can move the various tabs around for better comprehension

and comparison. Moreover, the initial comet metaphor had to be abandoned for the WebComets Touch implementation in order to better support the specific features of multi-touch interaction. This resulted in two visualizations that, while functionally almost identical, have a quite different representation of the data.

A major asset of WebComets visualization is the motif search interface which supports users in developing queries that focus both on content and context at the same time. This is even more relevant when motifs are coupled with the concept of property, as users could lock motifs to a ring and then take the TangibleRing with them, disallowing access to their query data. Also, for the purpose of searching structural data, most solutions still employ a descriptive language that needs to be entered manually. While more flexible, this approach offers no visual feedback on the structure that is being described, thus resulting in potential erroneous descriptions. On the other hand, building visual structural queries can also be problematic, if the representations do not support the entire vocabulary of the descriptive language.

In Chapter 4 we introduced the concepts of subjectivity measurements and visualization, focusing mostly on user emotional states and moments of insight. Our goal was to increase awareness of user emotions in both collaborative and individual work in order to positively impact effectiveness. One concern in this context is the detection of emotional states. More specifically, issues raised related to emotion detection were connected to the various levels of uncertainty. On the one hand, emotions are inherently subjective and volatile. As such, it is difficult to quantify to what extent the interpretations offered by the EEG system—or in fact any detection system—actually represents the emotional states of the users. One of the best approximations of user emotional states are represented by self-reports and databases of reported emotional states. Still, these approximations also have a subjective element that cannot yet be avoided with any approach. These aspects that are derived directly from the nature of emotions, should however not encourage us to disregard this dimension of user information. Instead, emotional reading should be regarded of affective hints guiding a collaboration or evaluation, and be inspected less like facts.

Furthermore, we noticed that EEG devices can represent a viable alternative for non-intrusiveness emotional readings. Still, one limitation we noticed is the inability to establish correlations between the detected affective states and the actions of the user, as these might not be the cause of the emotions.

The emotion visualization approaches in this thesis, while not complex to implement, represent concise solutions that are meant to easily integrate into different types of interfaces. Moreover, these representations convey emotions in an abstract way, by modulating visual attributes like color and shapes. However, some users have reported to prefer a more intuitive and natural visualizations, like emotional faces and avatars.

In terms of efficiency, while EmotionPrints and EmotionScent have been shown to convey the corresponding emotional information to the users, these solutions do have the downside of occupying additional screen space and adding complexity to the interface when incorporated in pre-existing visualizations. The result is an additional push towards information overload. This has been addressed by limiting the persistent nature of these visualizations and allowing users to control when they are represented on the interface. Furthermore, it should be stated that the emotion visualization approaches described in Chapter 4 are best suited to be used in interface design and evaluation as well as collaborative applications.

Besides raising the level of emotional awareness in collaborative visualization, in Chapter 5 we described a set of custom tangibles that address a number of requirements: reduced occlusion, increased precision, touch inside the object, combinability and cost efficiency. The contribution of TangibleRings in the field has also been confirmed by the inclusion of our approach in a recent survey paper on interactive lenses in visualization [282]. Still, the rings also have a set of limitations that need to be highlighted. TangibleRings have a non-variable diameter, making it problematic to nest two rings that are not of consecutive sizes. Furthermore, while the logical combinations of views and filters means that users should not require more than two rings at a time, the limited number of available diameters still represents restrictions. However, our evaluations and subsequent discussions have confirmed that users value the physical feedback obtained through tangibles. We hypothesize that future successful tangible designs will not only consider flexibility in function, but also the customization of the physical properties (e.g., size, shape, texture) of these objects.

Finally, the subjectivity evaluation of (collaborative) visualization systems offers a new dimension of insights into the dynamics of (collaborative) visual exploration. While subjectivity-based evaluations have been attempted with the help of other technologies (eye tracking, facial expressions), these approaches are less viable in co-located collaboration or visualization on large display due to their physical restrictions. Still, arguments can be made for a combined approach, were distinct modalities are employed in order to extract subjectivity information.

Characteristics of user-centered collaborative visualization In Chapter 1, a definition is given for the concept of user-centered collaborative visualization. Being the focal point of this thesis and a novel concept in the community, further information is required to consolidate the theory behind UCCV. Figure 6.1 highlights the interaction and awareness schematic behind UCCV. There are three types of interaction: social interaction between the users of the collaborative visualization, system interaction between the users and the visualization system, and action entity interaction between the actions each user executes on the system. In this context, an action entity (AE) is represented by all the actions a particular individual executes on the system using all the interaction modalities available to him/her. As a more symbolical explanation, action entities could also be defined as interaction avatars, i.e., the entity formed by all the interactions that take place between a user and the visualization.




Figure 6.1: Types of interaction and awareness in user-centered collaborative visualization. The system interaction is represented by the individual interactions executed upon the visualization by every user in a collaborative setting. At the same time, users influence each other through social interaction and communication, a level that requires social (e.g., emotional) awareness in order to ensure a coherent functioning of the group. At the other end, action entities (AE)—represented by the sum of activities of each user in the visualization—can influence each other, requiring goal, activity and change awareness for a coherent collaboration.

Moreover, based on our experience in designing and developing the technologies and approaches described in this thesis, we can formulate the following properties of UCCV:

1. The system interaction is not simply the sum of the individual interactions of every user. The effect a group of users has on a system is thus different from the effect each of the users would have executing the same actions separately.
2. The visualization system mirrors the social interaction within a group in the form of AE interaction. As users communicate and cooperate to achieve a goal, while at the same time addressing the specifics of their abilities and nature, similarly, AEs are composed of a set of actions that affect each other and are subject to their nature, modality and characteristics of the interactions supported by the visualization.
3. The environment of the UCCV can affect all three levels of interaction: social, system and active entity.
4. Awareness is vital on all levels of the collaboration process. Supporting data and situation awareness on the visualization side [82,103,260] should be complemented by the techniques that allow users to also increase their social awareness.

Social implications One tendency in terms of collaboration in visualization, and even in interactive systems in general, is represented by its move towards ubiquitous computing [219]—i.e., collaboration and interaction everywhere and anywhere. As computing-enabled collaboration will move to different aspects of our lives, the systems will also become more user-centered, focusing on the specifics of each user, each team, each particular experience and goal.

However, digital collaboration that is available wherever you are and exactly as you prefer it will also face a set of challenges. In a world where the digital systems are tuned to the needs and preferences of the users, the balance between user comfort and user goals might be upset, resulting in circumstances where task efficiency would have to suffer as (visualization) systems would try to accommodate the needs and desires of the users.

Perhaps an even greater implication of our research and UCCV in general is related to the realm of social interaction and communication. More specifically, user-centered approaches could reach new levels of depth in the following years, accessing personal and subjective information about users based on a large set of cues—e.g., subjectivity (emotion, mood, personality) measurements based on computer vision techniques, voice analysis, etc. In this context, the privacy and security of personal and intimate information becomes the focal point of the discussion.

It is clear that user-centered approaches need to obtain (objective and subjective) information about user personal aspects in order to support user-centered techniques. This however raises a couple of questions: is it acceptable to use techniques like computer vision algorithms to deduce and interpret user subjectivity in any environment without prior consent? Also, to what extent will individuals be willing to grant computing systems access to their subjective states?

In the first case, a thorough study will have to clarify to what extent these technologies would be allowed to collect subjective information, in what environments and to what purpose. Certainly, humans also have the ability to distinguish, for example based on facial expressions, if a person is happy or sad. However, a more systematic and perhaps even more precise and targeted reading of these subjective states raises a set of legal concerns that will have to be addressed at the right moment in time. Second, in terms of technology acceptance, it is unclear to what extent individuals would be willing to give up (part of) their privacy by allowing various computing systems to measure and interpret their subjective states. However, solutions that gather anonymized data, like the group affective tone, might represent a good compromise for considering this personal, subjective aspect of a group without overly intruding on the privacy of the users.

Collaborative sessions enabled by ubiquitous systems will also mean that these digital devices will become more integrated into our everyday lives, in terms of design and functionality, perhaps even hiding their computational nature. As a result, the dynamics of how we interact with these systems will be affected to the point where users will not make a distinction anymore between interacting and collaborating with everyday objects or digital systems.

Some of these implications (e.g., user privacy) have been partly addressed in this thesis. However, most of these considerations are still open for debate and beyond the scope of this research. Future studies and debates will certainly address issues related to the social implications of personal and subjectivity data, offering guiding principles for further research.

6.2 Conclusions

The overall objective of this thesis was to research and develop novel techniques for supporting a user-centered collaboration in the context of visualization systems. For this purpose, we started by highlighting a set of initial criteria centered around visualization metaphors, interaction techniques and user-specific dimensions in a collaborative environment (cf. Section 1.2).

- A. **Parallel browsing behavior** – Develop a novel visualization in order to test and evaluate our proposed UCCV techniques in a realistic (collaborative) visualization scenario.
- B. **Emotion visualization** – Create awareness of user subjective aspects in collaborative sessions in order to improve group cooperation, communication and overall social dynamics.
- C. **Tangible interaction** – Design custom interaction techniques that “naturally” support collaborative visualization and address issues like privacy, private and shared space, etc.
- D. **Subjectivity measures for evaluation** – Evaluate the proposed interaction and visualization techniques both through conventional and subjective measures in order to capture the quantitative and qualitative aspects of the proposed solutions.

Each of these individual criteria have been addressed as follows:

Parallel browsing behavior In Chapter 3 we highlight a novel interactive visualization for exploring and analyzing multi-session multi-user parallel browsing logs, called WebComets. The system is a novel time-series visualization of the tabbed navigation patterns of users, a visualization that supports advanced filtering and comparison of subgraphs through similarity measures and motif-based contextual search. A multi-touch version of this visualization tool is employed in Chapters 4 and 5 as the collaborative solution that allows us to test, evaluate and exemplify our user-centered techniques.

Emotion visualization A lightweight, commercial BCI headset is detailed in Chapter 4, where interpreted EEG data is employed to recognize and classify user emotional states. In Sections 4.3 and 4.4, a framework is described which enables the development of both multi-touch and desktop systems that support the incorporation of emotion visualization elements with the purpose of generating emotional awareness. The design considerations are highlighted, emphasizing the importance of interpreting and displaying emotional readings in real-time as well as at the display position where the interaction took place. We show that these visualizations, entitled EmotionPrints and EmotionScents, can be integrated in any setting on any multi-touch or desktop machine without requiring a complete redesign of the interface. Furthermore, a set of studies show that our affective representations increase emotional self-awareness, the awareness of other members of the team in a collaborative setting, and can support the evaluation of a system or interface.

Tangible interaction A novel concept for a passive tangible is presented in Chapter 5, called TangibleRings, that addresses some of the limitations presented by common tangibles (e.g., occlusion, lack of precision, no touch functionality on the surface of the tangible). TangibleRings can act similarly to virtual lenses in visualization, allowing users to generate their local, custom view or filter of the visualization and to manipulate it without affecting the other collaborators. Moreover, TangibleRings address issues like privacy and security, while at the same time supporting the construction of complex logical queries by combining and manipulating multiple rings.

The nesting of the concentric ring tangibles introduces a novel paradigm in touch interfaces, allowing users to manipulate and explore the dataset both locally—by manipulating multiple view parameters through the rings—and globally—by interacting with the rest of the tabletop surface.

Subjectivity measures for evaluation We validated the affective readings from the EEG headset against video logs/facial expressions and self-reports, both in simple and more complex visualization tasks. Our findings suggest that low-cost EEG headsets have the potential of offering reliable emotional readings for at least a set of emotional states. The subjectivity measurements have been employed as a means for qualitative evaluation of visualization systems in Section 5.2.3. The results of our study have shown that emotion recognition technology, in particularly more flexible and non-invasive approaches like wireless EEG, can offer valuable feedback about aspects of visualization design and interaction as well as insight into the particularities of a collaboration process.

In terms of the design guidelines for collaborative systems highlighted in [251], we believe we have managed to address the topic of supporting interpersonal interaction ([251]; Guideline 1) by capturing the group’s emotional state and increasing their awareness of each other’s subjectivity in the context of complex group dynamics. Furthermore, fluid transitions between activities ([251]; Guideline 2) as well as transitions between personal and group work ([251]; Guideline 3) have been facilitated through the physical and digital features of the TangibleRings, which enable users to not only manipulate and exchange rings, but also to focus on different types of tasks, with both local and global effects on the tabletop. By considering these aspects, our research introduces techniques that aid as well as build on the two main areas of experience that users have and need in the context of collaborating around tabletops: the experience of social interaction around a table and the one of positioning and manipulating objects on a surface.

Summary of contributions The scientific contributions of the research described in this thesis are:

1. Introduction of the novel concept of *user-centered collaborative visualization* (UCCV). Besides a definitions, specific value and properties of UCCV are further highlighted through the techniques described in Chapters 4 and 5 as well as through a set of specific characteristics of UCCV systems.
2. An interactive visualization that addresses the need for a tabbed browsing histories visualization, allowing the analysts to inspect the parallel browsing behavior of multiple users over different online sessions.
3. A visual motif builder that allows users to create patterns of navigation that incorporate both structural (navigation subgraph) and quantifiable (website attributes) information.
4. A set of evaluations inspecting the potential of portable commercial EEG headsets in the areas of user emotion and insight detection.
5. Three visualizations approaches for representing user and group affective states in various environments/interfaces (i.e., EmotionPrints, EmotionScents and GAT visualization).
6. A new design for passive tangibles for collaborative interaction on tabletop displays (i.e., TangibleRings). The associated interactions of the ring-shaped tangibles is tuned for collaboration and flexibility, allowing users to transition seamlessly between single-user views and manipulation of shared spaces.
7. The user subjectivity measures—i.e., emotional states and moments of insight—are employed in order to augment the quantitative evaluation of a collaborative visualization, offering a new dimension of information about the user experience in these settings.

6.3 Future Work

When devising the techniques and tools addressed in this thesis, one goal was to offer a complete view that includes sufficient details for reproduction and thorough evaluation of the results. At the same time, considering the inherent complexity of collaborative visualization, there is still room for many developments and extensions that can be applied to each of the tackled subjects in order to further explore both their validity and utility.

In terms of supporting the exploration and analysis of *user parallel browsing behavior*, we aim to incorporate automatic similarity search algorithms into the WebComets and WebComets Touch visualization environments that could support users by highlighting elements of high interest in the current query. We also plan to enhance the way websites are cataloged by including a dynamic semi-automatic updating of site categories. Furthermore, while we have shown in Chapter 3 that the WebComets approach is scalable in terms of the explored time frame, the number of branches defined by tabs and windows and the number of simultaneously inspected browsing logs, one important direction of our research in this area will be to further improve on scalability through new interaction metaphors. An example for such a metaphor would be a semantic zooming operator that changes the position of the loaded web pages on the temporal axis based on their relevance to the current query or analysis task.

A central aspect of this thesis is represented by the *subjectivity measurements and emotion visualization* for supporting collaborative visualization. In a first instance, we plan to evaluate and compare the potential of other devices and technologies capable of interpreting user emotions (e.g., electroencephalography (EEG), galvanic skin response (GSR), facial expression classification, eye tracking, etc.). These studies would extend the current one by further exploring the precision and level of intrusiveness of the devices, while at the same time being applied on a larger study group in order to reinforce the findings highlighted in this thesis.

Further, we aim to establish guidelines for subjectivity evaluations in the context of device-based affectivity detection and visualization systems. As highlighted in Sections 4.1 and 4.2, evaluations based on subjective aspects like emotional states and insights have the potential of capturing a dimension of the user experience in visualization that has been largely overlooked in this field.

We also plan to address some of the issues highlighted by the participants of our emotion visualization studies from Sections 4.3 and 4.4. Mainly, our intent is to improve the management of the emotion database and the privacy settings. To this end, but also to address the social implications of emotion visualization (see Section 6.1), we plan to implement a privacy management system that would span across all the emotion visualization techniques and scenarios, and that will enable users to clearly state the emotional readings, intervals and virtual objects for which they allow a representation to be computed. Further, we plan on enhancing the emotion visualization by combining the EEG measurements with data obtained from eye trackers in order to obtain a stronger correlation between the interface and the user experience, while at the same time removing some of the emotional data that is not related to the current task of the user.

Last but not least, in terms of *tangible interaction*, we plan to further improve the design of our TangibleRings in order to support more flexible and seamless interaction and collaboration. At the same time, we are working on a way to avoid the drawbacks of different ring diameters, while extending the interaction capabilities for the nested tangibles. At least as important are also the various application scenarios, where we plan to highlight the functionality of the TangibleRings by addressing questions related to medical visualization, as tangibles could improve navigating and filtering of medical and anatomic data in collaborative settings.

Bibliography

- [1] E. Adar, J. Teevan, and S. T. Dumais. Large scale analysis of web revisitation patterns. In *Proceedings of the 26th Annual SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, pages 1197–1206, New York, NY, USA, 2008. ACM Press.
- [2] M. Ahmadi and A. Erfanian. An on-line BCI system for hand movement control using real-time recurrent probabilistic neural network. In *Proceedings of the 4th International IEEE/EMBS Conference on Neural Engineering (NER '09)*, pages 367–370, 2009.
- [3] K. Aizawa, K. Ishijima, and M. Shiina. Summarizing wearable video. In *Proceedings of the IEEE International Conference on Image Processing (ICIP '01)*, volume 3, pages 398–401, 2001.
- [4] M. Akgun, G. K. Akilli, and K. Cagiltay. Bringing affect to human computer interaction. In *Affective Computing and Interaction*, pages 308–324. IGI Global, 2010.
- [5] E. W. Anderson. Evaluating visualization using cognitive measures. In *Proceedings of the BELIV Workshop: Beyond Time and Errors – Novel Evaluation Methods for Visualization (BELIV '12)*, pages 5:1–5:4, New York, NY, USA, 2012. ACM Press.
- [6] E. W. Anderson, K. C. Potter, L. E. Matzen, J. F. Shepherd, G. Preston, and C. T. Silva. A user study of visualization effectiveness using EEG and cognitive load. In H. Hauser, H. Pfister, and J. van Wijk, editors, *Proceedings of the 13th Eurographics/IEEE – VGTC Conference on Visualization (EuroVis '11)*, pages 791–800, Aire-la-Ville, Switzerland, Switzerland, 2011. Eurographics Association.

- [7] L. Applegate. Technology support for cooperative work: A framework for studying introduction and assimilation in organizations. *Journal Organizational Computing*, 1:11–39, 1991.
- [8] M. B. Arnold. *Emotion and Personality: Psychological aspects*, volume 1. Columbia University Press, New York, NY, USA, 1960.
- [9] P. M. Auble, J. J. Franks, and S. A. Soraci. Effort toward comprehension: Elaboration or “aha”? *Memory & Cognition*, 7(6):426–434, 1979.
- [10] A. Aula, N. Jhaveri, and M. Kaki. Information search and re-access strategies of experienced web users. *Proceedings of the 14th International Conference on World Wide Web (WWW '05)*, pages 583–592, 2005.
- [11] L. Aziz-Zadeh, J. Kaplan, and M. Iacoboni. “aha!”: The neural correlates of verbal insight solutions. *Human Brain Mapping*, 30(3):908–916, 2009.
- [12] N. A. Badcock, P. Mousikou, Y. Mahajan, P. de Lissa, J. Thie, and G. McArthur. Validation of the Emotiv EPOC EEG gaming system for measuring research quality auditory ERPs. *PeerJ*, 1(e38), 2013.
- [13] J. Bailenson, E. Pontikakis, I. Mauss, J. Gross, M. Jabon, C. Hutcherson, C. Nass, and O. John. Real-time classification of evoked emotions using facial feature tracking and physiological responses. *International Journal of Human-Computer Studies*, 66(5):303–317, 2008.
- [14] G. Bailly, A. Oulasvirta, T. Kötzing, and S. Hoppe. MenuOptimizer: Interactive optimization of menu systems. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*, pages 331–342, New York, NY, USA, 2013. ACM Press.
- [15] L. Bannon and S. Kjeld. CSCW: Four characters in search of a context. In *Proceedings of the 1st European Conference on Computer Supported Cooperative Work (ECSCW '89)*, pages 358–372, 1989.
- [16] G. Bateson. *Mind and nature: A necessary unity*. Dutton, Inc., New York, NY, USA, 1979.

- [17] P. Baudisch, T. Becker, and F. Rudeck. Lumino: Tangible blocks for tabletop computers based on glass fiber bundles. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*, pages 1165–1174, New York, NY, USA, 2010. ACM Press.
- [18] P. Baudisch, D. S. Tan, M. Collomb, D. C. Robbins, K. Hinckley, M. Agrawala, S. Zhao, and G. Ramos. Phosphor: explaining transitions in the user interface using afterglow effects. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology (UIST '06)*, pages 169–178. ACM Press, 2006.
- [19] S. Benbadis and D. Rielo. EEG artifacts. Medscape Reference, <http://emedicine.medscape.com/article/1140247-overview>, February 2014.
- [20] C. Berka, D. Levendowski, M. Lumicao, A. Yau, G. Davis, V. Zivkovic, R. Olmstead, P. Tremoulet, and P. Craven. EEG correlates of task engagement and mental workload in vigilance, learning, and memory tasks. *Aviation, Space, and Environmental Medicine*, 78(5):231–244, 2007.
- [21] E. Bertini, M. Rigamonti, and D. Lalanne. Extended excentric labeling. In *Proceedings of the 11th Eurographics / IEEE - VGTC Conference on Visualization (EuroVis '09)*, pages 927–934, Aire-la-Ville, Switzerland, 2009. Eurographics Association.
- [22] A. Bezerianos and P. Isenberg. Perception of visual variables on tiled wall-sized displays for information visualization applications. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2516–2525, 2012.
- [23] J. T. Biehl, M. Czerwinski, G. Smith, and G. G. Robertson. FASTDash: A visual dashboard for fostering awareness in software teams. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*, pages 1313–1322, New York, NY, USA, 2007. ACM.
- [24] E. A. Bier, M. C. Stone, K. Pier, W. Buxton, and T. D. DeRose. Toolglass and magic lenses: the see-through interface. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '93)*, pages 73–80, New York, NY, USA, 1993. ACM Press.

- [25] Y. Bouillon, F. Wendling, and F. Bartolomei. Computer-supported collaborative work (cscw) in biomedical signal visualization and processing. *IEEE Transactions on Information Technology in Biomedicine*, 3(1), 1999.
- [26] P. Bourne, M. Gribskov, G. Johnson, J. Moreland, S. Wavra, and H. Weisig. A prototype molecular interactive collaborative environment (mice). In *Proceedings of the Pacific Symposium on Biocomputing*, pages 118–129, 1998.
- [27] E. Bowden and M. Jung-Beeman. Aha! insight experience correlates with solution activation in the right hemisphere. *Psychonomic Bulletin and Review*, 10:730–737, 2003.
- [28] E. Bowden, M. Jung-Beeman, J. Fleck, and J. Kounios. New approaches to demystifying insight. *Trends in Cognitive Sciences*, 9(7):322–328, 2005.
- [29] G. Bower. How might emotions affect learning? In C. Svenåke, editor, *The handbook of emotion and memory: Research and theory*, pages 3–31. Lawrence Erlbaum Associates Ltd., 1992.
- [30] C. A. Brewer. Color use guidelines for data representation. In *Proceedings of the Section on Statistical Graphics, American Statistical Association, Alexandria VA*, pages 55–60, 1999.
- [31] I. Brewer, A. M. MacEachren, H. Abdo, J. Gundrum, and G. Otto. Collaborative geographic visualization: Enabling shared understanding of environmental processes. In *Proceedings of the IEEE Symposium on Information Visualization (InfoVis '00)*, pages 137–147, Washington, DC, USA, 2000. IEEE Computer Society.
- [32] M. Brewer and M. Hewstone. *Emotion and Motivation: Perspectives on Social Psychology*. Blackwell Publishing, 2004.
- [33] K. W. Brodlie, D. A. Duce, J. R. Gallop, J. P. R. B. Walton, and J. D. Wood. Distributed and collaborative visualization. *Computer Graphics Forum*, 23(2):223–251, 2004.
- [34] K. Bühler. Tatsachen und Probleme zu einer Psychologie der Denkvorgänge. *Archiv für Psychologie*, 9(315f):297–365, 2011.

- [35] N. Burns and S. Grove. The practice of nursing research conduct, critique, and utilization. *Philadelphia: W.B. Saunders and Co.*, 1997.
- [36] W. Büschel, U. Kister, M. Frisch, and R. Dachsel. T4 – transparent and translucent tangibles on tabletops. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '14)*, pages 81–88, New York, NY, USA, 2014. ACM Press.
- [37] G. L. Calhoun and G. R. McMillan. EEG-based control for human-computer interaction. In *Proceedings of the 3rd Symposium on Human Interaction with Complex Systems (HICS '96)*, pages 4–9, Washington, DC, USA, 1996. IEEE Computer Society.
- [38] A. Campbell, T. Choudhury, S. Hu, H. Lu, M. K. Mukerjee, M. Rabbi, and R. D. Raizada. NeuroPhone: brain-mobile phone interface using a wireless EEG headset. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Networking, Systems, and Applications on Mobile Handhelds (MobiHeld '10)*, pages 3–8, New York, NY, USA, 2010. ACM Press.
- [39] S. Card, J. Mackinlay, and B. Shneiderman. Readings in information visualization – using visualization to think. *Morgan Kaufmann*, 1999.
- [40] S. K. Card, G. G. Robertson, and W. York. The Webbook and the Web Forager: An information workspace for the World Wide Web. *Proceedings of the 14th Annual SIGCHI Conference on Human Factors in Computing Systems (CHI '96)*, pages 111–117, 1996.
- [41] S. Carpendale. Evaluating information visualizations. In A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, editors, *Information Visualization*, pages 19–45. Springer-Verlag, Berlin, Heidelberg, 2008.
- [42] S. Carpendale, A. Miede, T. Isenberg, S. D. Scott, R. Kruger, S. Habelski, U. Hinrichs, and K. Inkpen. Collaborative interaction on large tabletop displays. In *In Conference Supplement of the 2006 ACM Conference on Computer Supported Collaborative Work (CSCW '06)*, pages 57–58. ACM Press, 2006.
- [43] O. Casher and H. S. Rzepa. A chemical collaboratory using explorer EyeChem and the common client interface. *SIGGRAPH Computer Graphics*, 29(2):52–54, May 1995.

- [44] D. Cernea, A. Ebert, and A. Kerren. Detecting insight and emotion in visualization applications with a commercial EEG headset. In *Proceedings of the SIGRAD 2011 Conference on Evaluations of Graphics and Visualization – Efficiency, Usefulness, Accessibility, Usability*, pages 53–60, Stockholm, Sweden, 2011. Linköping University Electronic Press.
- [45] D. Cernea, A. Ebert, and A. Kerren. Visualizing group affective tone in collaborative scenarios. In *Proceedings of the Eurographics Conference on Visualization (EuroVis '14), Poster Abstract*, Swansea, Wales, UK, 2014.
- [46] D. Cernea, S. Mora, A. Perez, A. Ebert, A. Kerren, M. Divitini, D. G. de La Iglesia, and N. Otero. Tangible and wearable user interfaces for supporting collaboration among emergency workers. In *Proceedings of the 18th CRIWG Conference on Collaboration and Technology (CRIWG '12)*, volume 7493 of *LNCS*, pages 192–199, Duisburg, Germany, 2012. Springer.
- [47] D. Cernea, P.-S. Olech, A. Ebert, and A. Kerren. EEG-based measurement of subjective parameters in evaluations. In *Proceedings of the 14th International Conference on Human-Computer Interaction (HCII '11)*, volume 174 of *CCIS*, pages 279–283, Orlando, Florida, USA, 2011. Springer.
- [48] D. Cernea, P.-S. Olech, A. Ebert, and A. Kerren. Controlling in-vehicle systems with a commercial EEG headset: Performance and cognitive load. In *Visualization of Large and Unstructured Data Sets – Applications in Geospatial Planning, Modeling and Engineering (VLUDS '12)*, pages 113–122, Schloss Dagstuhl, Leibniz-Zentrum für Informatik, 2012. OpenAccess Series in Informatics (OASIS).
- [49] D. Cernea, P.-S. Olech, A. Ebert, and A. Kerren. Measuring subjectivity – supporting evaluations with the Emotiv EPOC neuroheadset. *Künstliche Intelligenz – KI Journal, Special Issue on Human-Computer Interaction*, 26(2):177–182, 2012.
- [50] D. Cernea, I. Truderung, A. Kerren, and A. Ebert. WebComets: A tab-oriented approach for browser history visualization. In *Proceedings of the International Conference on Information Visualization Theory and Applications (IVAPP '13)*, pages 439–450, Barcelona, Spain, 2013. SciTePress.

- [51] D. Cernea, I. Truderung, A. Kerren, and A. Ebert. An interactive visualization for tabbed browsing behavior analysis. *Computer Vision, Imaging and Computer Graphics – Theory and Applications, Communications in Computer and Information Science*, 458:1–16, 2014.
- [52] D. Cernea, C. Weber, A. Ebert, and A. Kerren. Emotion Scents – a method of representing user emotions on GUI widgets. In *Proceedings of the SPIE 2013 Conference on Visualization and Data Analysis (VDA '13)*, volume 8654, Burlingame, CA, USA, 2013. IS&T/SPIE.
- [53] D. Cernea, C. Weber, A. Ebert, and A. Kerren. Emotion-prints: Interaction-driven emotion visualization on multi-touch interfaces. In *Proceedings of the SPIE 2015 Conference on Visualization and Data Analysis (VDA '15)*, volume 9397, Burlingame, CA, USA, 2015. IS&T/SPIE (to be published).
- [54] L. Chan, S. Müller, A. Roudaut, and P. Baudisch. CapStones and ZebraWidgets: sensing stacks of building blocks, dials and sliders on capacitive touch screens. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*, pages 2189–2192, New York, NY, USA, 2012. ACM Press.
- [55] R. Chang, C. Ziemkiewicz, T. Green, and W. Ribarsky. Defining insight for visual analytics. *IEEE Computer Graphics and Applications*, 29(2):14–17, 2009.
- [56] S. Chawla and P. Bedi. Improving information retrieval precision by finding related queries with similar information need using information scent. In *Proceedings of the 1st International Conference on Emerging Trends in Engineering and Technology (ICETET '08)*, pages 486–491, Washington, DC, USA, 2008. IEEE Computer Society.
- [57] I. Chtivelband. *HistoryLane: Web Browser History Visualization Method*. M.sc. thesis, University of Kaiserslautern, Germany and Blekinge Institute of Technology, Sweden, 2012.
- [58] A. Cockburn, S. Greenberg, S. Jones, B. McKenzie, and M. Moyle. Improving web page revisitation: Analysis, design and evaluation. *IT & Society*, 1(3):159–183, 2003.
- [59] D. Cogburn. HCI in the so-called developing world: What’s in it for everyone. *Interactions*, 10(2):80–87, 2003.

- [60] B. Cooper, P. Brna, and A. Martins. Effective affective in intelligent systems – building on evidence of empathy in teaching and learning. *Affective Interactions: Towards a New Generation of Computer Interfaces*, pages 21–34, 2000.
- [61] N. Couture, G. Riviere, and P. Reuter. GeoTUI: A tangible user interface for geoscience. In *Proceedings of the 2nd International Conference on Tangible and Embedded Interaction (TEI '08)*, pages 89–96, Bonn, Germany, 2008. ACM Press.
- [62] J. Cugini and J. Scholtz. VISVIP: 3D visualization of paths through websites. In *Proceedings of International Workshop on Web-Based Information Visualization (WebVis '99)*, pages 259–263, Florence, Italy, 1999.
- [63] A. Cummings. Contextual characteristics and employee creativity: Affect at work. *Proceedings of the 13th Annual Conference, Society for Industrial Organizational Psychology*, 1998.
- [64] A. Damasio. *Descarte's Error: Emotion, Reason and the Human Brain*. Avon, New York, NY, USA, 1994.
- [65] L. Damianos, L. Hirschman, R. Kozierok, J. Kurtz, A. Greenberg, K. Walls, S. Laskowski, and J. Scholtz. Evaluation for collaborative systems. *ACM Computing Surveys*, 31(2es), June 1999.
- [66] R. Davidson. Affective neuroscience and psychophysiology: Toward a synthesis. *Psychophysiology*, 40:655–665, 2003.
- [67] R. Davidson. What does the prefrontal cortex do in affect: perspectives on frontal eeg asymmetry research. *Biological Psychology*, 67(1–2):219–233, 2004.
- [68] R. J. Davidson. Affective style and affective disorders: Perspectives from affective neuroscience. *Cognition and Emotion*, 12(3):307–330, 1998.
- [69] P. Desmet. Measuring emotion: Development and application of an instrument to measure emotional responses to products. In M. A. Blythe, K. Overbeeke, A. F. Monk, and P. C. Wright, editors, *Funology*, pages 111–123. Kluwer Academic Publishers, Norwell, MA, USA, 2004.

- [70] S. D’Mello and R. A. Calvo. Beyond the basic emotions: What should affective computing compute? In *Extended Abstracts on Human Factors in Computing Systems (CHI EA ’13)*, pages 2287–2294, New York, NY, USA, 2013. ACM Press.
- [71] P. Doemel. Webmap – a graphical hypertext navigation tool. In *Proceedings of the 2nd International Conference on World Wide Web (WWW ’95)*, volume 28, pages 85–97, 1995.
- [72] C. D. Dryer. Dominance and valence: A two-factor model for emotion in HCI. In *Emotional and Intelligent: The Tangled Knot of Cognition*, TR FS-98-03, pages 75–81, Menlo Park, CA, USA, 1998. AAAI Press.
- [73] P. Dubroy and R. Balakrishnan. A study of tabbed browsing among mozilla firefox users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’10)*, pages 673–682, New York, NY, USA, 2010. ACM Press.
- [74] G. W. Dunnington. *Carl Friedrich Gauss: Titan of Science*. Mathematical Association of America, 2004.
- [75] M. Duvinage, T. Castermans, M. Petieau, T. Hoellinger, G. Cheron, and T. Dutoit. Performance of the Emotiv EPOC headset for P300-based applications. *BioMedical Engineering OnLine*, 12(1):1–15, 2013.
- [76] A. Ebert, C. Weber, D. Cernea, and S. Petsch. TangibleRings: Nestable circular tangibles. In *Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI ’13)*, pages 1617–1622, Paris, France, 2013.
- [77] S. G. Eick. Visual analysis of website browsing patterns. In *Visual Interfaces to Digital Libraries, JCDL 2002 Workshop*, pages 65–80, London, UK, 2002. Springer-Verlag.
- [78] P. Ekman. Universals and cultural differences in facial expressions of emotion. *Nebraska Symposium on Motivation 1971*, 19:207–283, 1972.
- [79] P. Ekman. An argument for basic emotions. *Cognition and Emotion*, 6(3/4):160–200, 1992.

- [80] M. S. El-Nasr, J. Morie, and A. Drachen. A scientific look at the design of aesthetically and emotionally engaging interactive entertainment experiences. In *Affective Computing and Interaction: Psychological, Cognitive and Neuroscientific Perspectives*, Information Science Publishing, IGI Global, pages 281–307, 2010.
- [81] N. Elmqvist, A. Vande Moere, H.-C. Jetter, D. Cernea, H. Reiterer, and T. J. Jankun-Kelly. Fluid interaction for information visualization. *Information Visualization Journal (IVS), Special Issue on Information Visualization: State of the Field and New Research Directions*, 10(4):327–340, 2011.
- [82] M. R. Endsley. *Designing for Situation Awareness: An Approach to User-Centered Design*. CRC Press, Inc., Boca Raton, FL, USA, 2 edition, 2011.
- [83] M. Eric, E. Z. Ayers, and J. T. Stasko. Using graphic history in browsing the World Wide Web. *Proceedings of the 2nd International Conference on World Wide Web (WWW '95)*, 28:1–7, 1995.
- [84] J.-D. Fekete, J. J. Wijk, J. T. Stasko, and C. North. The value of information visualization. In A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, editors, *Information Visualization*, pages 1–18. Springer-Verlag, Berlin, Heidelberg, 2008.
- [85] G. W. Fitzmaurice. Bricks: Laying the foundations for graspable user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '95)*, pages 442–449, New York, NY, USA, 1995. ACM Press.
- [86] C. Forlines, A. Esenther, C. Shen, D. Wigdor, and K. Ryall. Multi-user, multi-display interaction with a single-user, single-display geospatial application. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology (UIST '06)*, pages 273–276, New York, NY, USA, 2006. ACM.
- [87] C. Forlines and R. Lilien. Adapting a single-user, single-display molecular visualization application for use in a multi-user, multi-display environment. In *Proceedings of the Working Conference on Advanced Visual Interfaces (AVI '08)*, pages 367–371, New York, NY, USA, 2008. ACM.

- [88] C. Forlines and C. Shen. DTLens: Multi-user tabletop spatial data exploration. In *Proceedings of the 18th Annual ACM Symposium on User Interface Software and Technology (UIST '05)*, pages 119–122, New York, NY, USA, 2005. ACM Press.
- [89] B. Fredrickson. The value of positive emotions. *American Scientist*, 91:330–335, 2003.
- [90] N. Frijda. Mood. *The Oxford companion to emotion and the affective sciences*, pages 258–259, 2009.
- [91] R. Gandhi, K. Girish, B. B. Bederson, and B. Shneiderman. Domain name based visualization of web histories in a zoomable user interface. In *Proceedings of 11th International Workshop on Database and Expert Systems Applications (DEXA '00)*, pages 591–598, 2000.
- [92] O. García, J. Favela, G. Licea, and R. Machorro. Extending a collaborative architecture to support emotional awareness. In *Proceedings of the Workshop on Emotion-Based Agent Architectures, Autonomous Agents (EBAA '99)*, pages 46–52, 1999.
- [93] O. García, J. Favela, and R. Machorro. Emotional awareness in collaborative systems. In *Proceedings of the String Processing and Information Retrieval Symposium & International Workshop on Groupware (SPIRE '99)*, pages 296–303, Washington, DC, USA, 1999. IEEE Computer Society.
- [94] J. George. Personality, affect, and behavior in groups. *Journal of Applied Psychology*, 75:107–116, 1990.
- [95] J. George. Group affective tone. *Handbook of work group psychology*, 75:77–93, 1996.
- [96] D. Gökçay. Emotional axes: Psychology, psychophysiology and neuroanatomical correlates. In D. Gökçay and G. Yildirim, editors, *Affective Computing and Interaction: Psychological, Cognitive and Neuroscientific Perspectives*, pages 56–73. IGI Global, 2011.
- [97] L. Goldberg. An alternative “description of personality”: The big-five factor structure. *Journal of Personality and Social Psychology*, 59(6):1216–1229, 1990.

- [98] E. J. Gomez, F. del Pozo, E. J. Ortiz, N. Malpica, and H. Rahms. A broadband multimedia collaborative system for advanced teleradiology and medical imaging diagnosis. *IEEE Transactions on Information Technology in Biomedicine*, 2(3):146–155, September 1998.
- [99] D. Grimes, D. S. Tan, S. E. Hudson, P. Shenoy, and R. P. Rao. Feasibility and pragmatics of classifying working memory load with an electroencephalograph. In *Proceeding of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, pages 835–844, New York, NY, USA, 2008. ACM Press.
- [100] I. J. Grimstead, D. W. Walker, and N. J. Avis. Collaborative visualization: A review and taxonomy. In *Proceedings of the 9th IEEE International Symposium on Distributed Simulation and Real-Time Applications (DS-RT '05)*, pages 61–69, Washington, DC, USA, 2005. IEEE Computer Society.
- [101] E. F. Gross. Adolescent internet use: What we expect, what teens report. *Applied Developmental Psychology*, 25:633–649, 2004.
- [102] L. Harrison, D. Skau, S. Franconeri, A. Lu, and R. Chang. Influencing visual judgment through affective priming. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, pages 2949–2958, New York, NY, USA, 2013. ACM Press.
- [103] J. Heer and M. Agrawala. Design considerations for collaborative visual analytics. *Information Visualization*, 7(1):49–62, March 2008.
- [104] J. Heer, F. Ham, S. Carpendale, C. Weaver, and P. Isenberg. Creation and collaboration: Engaging new audiences for information visualization. In A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, editors, *Information Visualization*, pages 92–133. Springer-Verlag, Berlin, Heidelberg, 2008.
- [105] J. Heer, J. Mackinlay, C. Stolte, and M. Agrawala. Graphical histories for visualization: Supporting analysis, communication, and evaluation. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1189–1196, November 2008.

- [106] D. Heger, F. Putze, and T. Schultz. Online recognition of facial actions for natural EEG-based BCI applications. In *Proceedings of the 4th International Conference on Affective Computing and Intelligent Interaction (ACII '11)*, volume 2, pages 436–446, Berlin, Heidelberg, Germany, 2011. Springer-Verlag.
- [107] S. H. Hemenover. Individual differences in rate of affect change: Studies in affective chronometry. *Journal of Personality and Social Psychology*, 85(1):121–131, 2003.
- [108] B. Hengeveld, C. Hummels, K. Overbeeke, R. Voort, H. van Balkom, and J. de Moor. Let me actuate you. In *Proceedings of the 2nd International Conference on Tangible and Embedded Interaction (TEI '08)*, pages 159–166, New York, NY, USA, 2008. ACM Press.
- [109] P. Henline. Eight collaboratory summaries. *Interactions*, 5(3):66–72, 1998.
- [110] F. Hennecke, F. Berwein, and A. Butz. Optical pressure sensing for tangible user interfaces. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '11)*, pages 45–48, New York, NY, USA, 2011. ACM Press.
- [111] R. R. Hightower, L. T. Ring, J. I. Helfman, B. B. Bederson, and J. D. Hollan. Graphical multiscale web histories: A study of PadPrints. In *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology (UIST '98)*, pages 58–65, New York, NY, USA, 1998. ACM Press.
- [112] J. Hill and C. Gutwin. Awareness support in a groupware widget toolkit. In *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work (GROUP '03)*, pages 258–267, New York, NY, USA, 2003. ACM Press.
- [113] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell. Passive real-world interface props for neurosurgical visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Celebrating interdependence (CHI '94)*, pages 452–458, New York, NY, USA, 1994. ACM Press.

- [114] L. Hole and N. Millard. Affective Widgets: Supporting the daily work experiences of contact centre advisors. In *Proceedings of the User Experience Workshop at the 4th Nordic Conference on Human-Computer Interaction (NordiCHI '06)*, 2006.
- [115] K. Hone. Empathic agents to reduce user frustration: The effects of varying agent characteristics. *Interacting with Computers*, 18(2):227–245, 2006.
- [116] R. Horlings, D. Datcu, and L. J. M. Rothkrantz. Emotion recognition using brain activity. In *Proceedings of the 9th International Conference on Computer Systems and Technologies and Workshop for PhD Students in Computing (CompSysTech '08)*, pages II.1–1, New York, NY, USA, 2008. ACM Press.
- [117] A. H. Huang, D. C. Yen, and X. Zhang. Exploring the potential effects of emoticons. *Information and Management*, 45(7):466–473, November 2008.
- [118] C.-J. Huang, E. Y.-L. Do, and M. D. Gross. MouseHaus Table: A physical interface for urban design. In *Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology (UIST '03)*, Vancouver, BC, Canada, 2003. ACM Press.
- [119] J. Huang and R. W. White. Parallel browsing behavior on the web. In *Proceedings of the 21st ACM Conference on Hypertext and Hypermedia (HT '10)*, pages 13–18, 2010.
- [120] S. Huang. Iconic history. <http://shan-huang.com/browserdatavis/>, February 2014.
- [121] G. Huisman, M. van Hout, E. van Dijk, T. van der Geest, and D. Heylen. LEMtool: measuring emotions in visual interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, pages 351–360, New York, NY, USA, 2013. ACM Press.
- [122] P. S. Inventado, R. S. Legaspi, M. Numao, and M. Suarez. Observatory: A tool for recording, annotating and reviewing emotion-related data. In *Proceedings of the 3rd International Conference on Knowledge and Systems Engineering (KSE '11)*, pages 261–265. IEEE, 2011.
- [123] A. Isen, K. Daubman, and G. Nowicki. Positive affect facilitates creative problem solving. *Journal of Personality and Social Psychology*, 52:1122–1131, 1987.

- [124] P. Isenberg, N. Elmqvist, J. Scholtz, D. Cernea, K.-L. Ma, and H. Hagen. Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization Journal (IVS), Special Issue on Information Visualization: State of the Field and New Research Directions*, 10(4):310–326, 2011.
- [125] P. Isenberg and D. Fisher. Collaborative brushing and linking for co-located visual analytics of document collections. In *Proceedings of the 11th Eurographics / IEEE – VGTC Conference on Visualization (EuroVis '09)*, pages 1031–1038, Aire-la-Ville, Switzerland, Switzerland, 2009. Eurographics Association.
- [126] P. Isenberg, U. Hinrichs, M. Hancock, and S. Carpendale. Digital tables for collaborative information exploration. In C. Müller-Tomfelde, editor, *Tabletops – Horizontal Interactive Displays*, Human-Computer Interaction Series, pages 387–405. Springer London, 2010.
- [127] H. Ishii. The tangible user interface and its evolution. *Communications of the ACM*, 51(6):32–36, June 2008.
- [128] H. Ishii and B. Ullmer. Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '97)*, pages 234–241, New York, NY, USA, 1997. ACM Press.
- [129] N. Jatupaiboon, S. Pan-ngum, and P. Israsena. Emotion classification using minimal EEG channels and frequency bands. In *Proceedings of the 10th International Joint Conference on Computer Science and Software Engineering (JCSSE '13)*, pages 21–24, May 2013.
- [130] H.-C. Jetter, J. Gerken, M. Zöllner, H. Reiterer, and N. Milic-Frayling. Materializing the query with facet-streams: a hybrid surface for collaborative search on tabletops. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*, pages 3013–3022. ACM Press, 2011.
- [131] S. Jordà, G. Geiger, M. Alonso, and M. Kaltenbrunner. The reacTable: exploring the synergy between live music performance and tabletop tangible interfaces. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction (TEI '07)*, pages 139–146, New York, NY, USA, 2007. ACM Press.

- [132] T.-P. Jung, S. Makeig, C. Humphries, T.-W. Lee, M. J. Mckeown, V. Iragui, and T. J. Sejnowski. Removing electroencephalographic artifacts by blind source separation. *Psychophysiology*, 37(2):163–178, 2000.
- [133] M. Jung-Beeman, E. Bowden, J. Haberman, J. Frymiare, S. Arambel-Liu, R. Greenblatt, P. Reber, and J. Kounios. Neural activity when people solve verbal problems with insight. *PLoS Biology*, 2(4): e97, 2004.
- [134] S. Kaasten and S. Greenberg. Integrating back, history and bookmarks in web browsers. In *Extended Abstracts of the ACM Conference of Human Factors in Computing Systems (CHI '01)*, pages 1–2. ACM Press, 2000.
- [135] S. Kaasten, S. Greenberg, and C. Edwards. How people recognize previously seen web pages from titles, URLs and thumbnails. In X. Faulkner, J. Finlay, and F. Détienne, editors, *People and Computers XVI – Memorable Yet Invisible*, pages 247–265. Springer London, 2002.
- [136] Y. Kakehi, M. Iida, and T. Naemura. Tablescape Plus: Upstanding tiny displays on tabletop display. In *ACM SIGGRAPH 2006 Sketches*, SIGGRAPH '06, New York, NY, USA, 2006. ACM Press.
- [137] D. P. Käser, M. Agrawala, and M. Pauly. FingerGlass: Efficient multiscale interaction on multitouch screens. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*, pages 1601–1610. ACM Press, 2011.
- [138] J. Katz and P. Aspden. Motivations for and barriers to internet usage: results of a national public opinion survey. *Internet Research: Electronic Networking Applications and Policy*, 7:170–188, 1997.
- [139] M. Kellar, C. Watters, and K. M. Inkpen. An exploration of web-based monitoring: implications for design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*, pages 377–386, New York, NY, USA, 2007. ACM Press.
- [140] A. Kerren, A. Ebert, and J. Meyer, editors. *Human-Centered Visualization Environments*, volume 4417 of *LNCS*. Springer-Verlag, Heidelberg, Germany, 2007.

- [141] K. Kim, W. Javed, C. Williams, N. Elmqvist, and P. Irani. Hugin: A framework for awareness and coordination in mixed-presence collaborative information visualization. In *ACM International Conference on Interactive Tabletops and Surfaces (ITS '10)*, pages 231–240, New York, NY, USA, 2010. ACM.
- [142] M. Kipp. Anvil – a generic annotation tool for multimodal dialogue. In *Proceedings of the 7th European Conference on Speech Communication and Technology (Eurospeech)*, pages 1367–1370, 2001.
- [143] J. Klein, Y. Moon, and R. Picard. This computer responds to user frustration: theory, design and results. *Interacting with Computers*, 14(2):119–140, 2002.
- [144] S. R. Klemmer, M. W. Newman, R. Farrell, M. Bilezikjian, and J. A. Landay. The designers’ outpost: A tangible interface for collaborative web site. In *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology (UIST '01)*, pages 1–10, New York, NY, USA, 2001. ACM.
- [145] S. Klum, P. Isenberg, R. Langner, J.-D. Fekete, and R. Dachsel. Stackables: combining tangibles for faceted browsing. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '12)*, pages 241–248. ACM Press, 2012.
- [146] M.-J. Ko, G.-H. Oh, and K.-W. Bae. A study on new gameplay experience based on brain-computer interface. *Journal of Korea Game Society*, 9:31–44, December 2009.
- [147] H. Koike, W. Nishikawa, and K. Fukuchi. Transparent 2-D markers on an LCD tabletop system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*, pages 163–172, New York, NY, USA, 2009. ACM Press.
- [148] V. Konovalov and I. Serikov. Characteristics of the galvanic skin response and electrocardiogram in active and passive subjects under test conditions. *Human Physiology*, 32:578–583, 2006.
- [149] C. A. Kothe, S. Makeig, and J. A. Onton. Emotion recognition from EEG during self-paced emotional imagery. In *Proceedings of the 5th International Conference on Affective Computing and Intelligent Interaction (ACII '13)*, pages 855–858, 2013.

- [150] J. Kounios, J. Fleck, D. Green, L. Payne, J. Stevenson, E. Bowden, and M. Jung-Beeman. The origins of insight in resting-state brain activity. *Neuropsychologia*, 46:281–291, 2008.
- [151] J. Kounios, J. L. Frymiare, E. M. Bowden, J. I. Fleck, K. Subramaniam, T. B. Parrish, and M. Jung-Beeman. The prepared mind: Neural activity prior to problem presentation predicts subsequent solution by sudden insight. *Psychological Science*, 17(10):882–890, 2006.
- [152] U. Krčadinac, J. Jovanović, and V. Devedžić. Visualizing the affective structure of students interaction. In S. K. Cheung, J. Fong, L.-F. Kwok, K. Li, and R. Kwan, editors, *Hybrid Learning*, volume 7411 of *LNCS*, pages 23–34. Springer Berlin Heidelberg, 2012.
- [153] M. Krstajic, E. Bertini, and D. Keim. CloudLines: Compact display of event episodes in multiple time-series. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2432–2439, December 2011.
- [154] R. Kruger, S. Carpendale, S. D. Scott, and S. Greenberg. Roles of orientation in tabletop collaboration: Comprehension, coordination and communication. *Computer Supported Cooperative Work*, 13(5–6):501–537, Dec. 2004.
- [155] M. Labaj and M. Bieliková. Tabbed browsing behavior as a source for user modeling. In S. Carberry, S. Weibelzahl, A. Micarelli, and G. Semeraro, editors, *User Modeling, Adaptation, and Personalization*, volume 7899 of *LNCS*, pages 388–391. Springer Berlin Heidelberg, 2013.
- [156] R. Larsen and E. Diener. Promises and problems with the circumplex model of emotion. *Review of personality and social psychology*, 13:25–59, 1992.
- [157] R. Leeb, D. Friedman, G. R. Müller-Putz, R. Scherer, M. Slater, and G. Pfurtscheller. Self-paced (asynchronous) BCI control of a wheelchair in virtual environments: a case study with a tetraplegic. *Computational Intelligence and Neuroscience – Brain-Computer Interfaces: Towards Practical Implementations and Potential Applications*, 2007:7:1–7:12, April 2007.
- [158] J. Lehrer. The eureka hunt: Why do good ideas come to us when they do? *The New Yorker*, pages 40–45, July 28, 2008.

- [159] L. A. Liikkanen, G. Jacucci, and M. Helin. ElectroEmotion – a tool for producing emotional corpora collaboratively. In *Proceedings of the 3rd International Conference on Affective Computing and Intelligent Interaction (ACII '09)*, pages 1–7, 2009.
- [160] L. A. Liikkanen, G. Jacucci, E. Huvio, T. Laitinen, and E. Andre. Exploring emotions and multimodality in digitally augmented puppeteering. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '08)*, pages 339–342. ACM Press, 2008.
- [161] R. Likert. A technique for the measurement of attitudes. *Archives of Psychology*, 140:1–55, 1932.
- [162] C. L. Lisetti and F. Nasoz. MAUI: A multimodal affective user interface. In *Proceedings of the 10th ACM International Conference on Multimedia (MULTIMEDIA '02)*, pages 161–170, New York, NY, USA, 2002. ACM Press.
- [163] H. Liu, T. Selker, and H. Lieberman. Visualizing the affective structure of a text document. In *Extended Abstracts on Human Factors in Computing Systems (CHI EA '03)*, pages 740–741, New York, NY, USA, 2003. ACM Press.
- [164] Y. Liu, O. Sourina, and M. K. Nguyen. Real-time EEG-based human emotion recognition and visualization. In *Proceedings of the 2010 International Conference on Cyberworlds (CW '10)*, pages 262–269, Washington, DC, USA, 2010. IEEE Computer Society.
- [165] A. Lu, L. Harrison, and R. Chang. Exploring the impact of emotion on visual judgement. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST '12)*, pages 227–228, Washington, DC, USA, 2012. IEEE Computer Society.
- [166] A. Luneski and R. K. Moore. Affective computing and collaborative networks: Towards emotion-aware interaction. In L. M. Camarinha-Matos and W. Picard, editors, *Pervasive Collaborative Networks*, volume 283 of *IFIP International Federation for Information Processing*, pages 315–322. Springer US, 2008.
- [167] J. Luo and K. Niki. Function of hippocampus in “insight” of problem solving. *Hippocampus*, 13:316–323, 2003.

- [168] A. MacEachren and J. Ganter. A pattern identification approach to cartographic visualization. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 27(2):64–81, 1990.
- [169] A. M. MacEachren, M. Gahegan, W. Pike, I. Brewer, G. Cai, E. Lengerich, and F. Hardisty. Geovisualization for knowledge construction and decision support. *IEEE Computer Graphics and Applications*, 24(1):13–17, January 2004.
- [170] E. Maguire, P. Rocca-Serra, S.-A. Sansone, J. Davies, and M. Chen. Visual compression of workflow visualizations with automated detection of macro motifs. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2576–2585, 2013.
- [171] W. Mähr. emotion: An outline of an emotion aware self-adaptive user interface. In *Proceedings of the 2nd Scandinavian Student Interaction Design Research Conference (SIDER '06)*, pages 25–26. Gothenburg, Sweden, 2006.
- [172] N. Mahyar, A. Sarvghad, and M. Tory. Note-taking in co-located collaborative visual analytics: Analysis of an observational study. *Information Visualization*, 11(3):190–204, July 2012.
- [173] I. H. Manssour, C. Maria, and D. S. Freitas. Collaborative visualization in medicine. In *Proceedings of the International Conference in Central Europe on Computer Graphic, Visualization and Interactive Digital Media (WSCG '00)*, pages 266–273, 2000.
- [174] V. Maquil, T. Psik, and I. Wagner. The ColorTable: A design story. In *Proceedings of the 2nd International Conference on Tangible and Embedded Interaction (TEI '08)*, pages 97–104, New York, NY, USA, 2008. ACM Press.
- [175] V. Maquil, E. Ras, and O. Zephir. Understanding the characteristics of metaphors in tangible user interfaces. In M. Eibl and M. Ritter, editors, *Proceedings of the Tagung Mensch & Computer Workshop*, pages 47–52, Chemnitz, Germany, 2011. Universitätsverlag Chemnitz.
- [176] F. T. Marchese and N. Brajkovska. Fostering asynchronous collaborative visualization. In *Proceedings of the 11th International Conference Information Visualization (IV '07)*, pages 185–190, Washington, DC, USA, 2007. IEEE Computer Society.

- [177] J. Marco, E. Cerezo, S. Baldassarri, E. Mazzone, and J. C. Read. Bringing tabletop technologies to kindergarten children. In *Proceedings of the 23rd British HCI Group Annual Conference on People and Computers: Celebrating People and Technology (BCS-HCI '09)*, pages 103–111, Swinton, UK, 2009. British Computer Society.
- [178] H. Margolis. *Patterns, Thinking, and Cognition: A Theory of Judgment*. University Of Chicago Press, 1 edition, 1990.
- [179] I. Martinovic, D. Davies, M. Frank, D. Perito, T. Ros, and D. Song. On the feasibility of side-channel attacks with brain-computer interfaces. In *Proceedings of the 21st USENIX Conference on Security Symposium (Security'12)*, pages 34–34, Berkeley, CA, USA, 2012. USENIX Association.
- [180] M. Masuch, K. Hartman, and G. Schuster. Emotional agents for interactive environments. In *Proceedings of the 4th International Conference on Creating, Connecting and Collaborating through Computing (C5 '06)*, pages 96–102, Washington, DC, USA, 2006. IEEE Computer Society.
- [181] Y. Y. Mathieu. Annotation of emotions and feelings in texts. In *Proceedings of the 1st International Conference on Affective Computing and Intelligent Interaction (ACII '05)*, pages 350–357, Berlin, Heidelberg, 2005. Springer-Verlag.
- [182] J. D. Mayer and P. Salovey. What is emotional intelligence? In P. Salovey and D. Sluyter, editors, *Emotional development and emotional intelligence: Implications for educators*, pages 3–31. New York: Basic Books, 1997.
- [183] R. Mazza. *Introduction to Information Visualization*. Springer Publishing Company, Incorporated, 1 edition, 2009. Chapter 6.2.
- [184] D. McDuff, A. Karlson, A. Kapoor, A. Roseway, and M. Czerwinski. AffectAura: An intelligent system for emotional memory. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*, pages 849–858, New York, NY, USA, 2012. ACM Press.
- [185] A. Mehrabian. Framework for a comprehensive description and measurement of emotional states. *Genetic, social, and general psychology monographs*, 121(3):339–361, August 1995.

- [186] B. Mesquita and N. Frijda. Cultural variations in emotions: A review. *Psychological Bulletin*, 112(2):179–204, 1992.
- [187] P. Michel and R. El Kaliouby. Real time facial expression recognition in video using support vector machines. In *Proceedings of the 5th International Conference on Multimodal Interfaces (ICMI '03)*, pages 258–264, New York, NY, USA, 2003. ACM Press.
- [188] M. Mikhail, K. El-Ayat, R. El Kaliouby, J. Coan, and J. J. B. Allen. Emotion detection using noisy EEG data. In *Proceedings of the 1st Augmented Human International Conference (AH '10)*, pages 1–7, New York, NY, USA, 2010. ACM Press.
- [189] N. Millard and L. Hole. In the moodie: Using 'affective widgets' to help contact centre advisors fight stress. In C. Peter and R. Beale, editors, *Affect and Emotion in Human-Computer Interaction*, volume 4868 of *LNCS*, pages 186–193. Springer, 2008.
- [190] Y. Miyata and D. Norman. Psychological issues in support of multiple activities. *User Centered System Design*, pages 265–284, 1986.
- [191] G. G. Molina, T. Tsoneva, and A. Nijholt. Emotional brain-computer interfaces. In J. Cohn, A. Nijholt, and M. Pantic, editors, *Proceedings of the 3rd International Conference on Affective Computing and Intelligent Interaction (ACII '09)*, pages 138–146, Los Alamitos, September 2009. IEEE Computer Society Press.
- [192] Y. E. Moon. *Similarity Effects in Human-computer Interaction: Effects of User Personality, Computer Personality, and User Control on Attraction and Attributions of Responsibility*. PhD thesis, Stanford University, Stanford, CA, USA, 1996. AAI9630362.
- [193] M. R. Morris, K. Ryall, C. Shen, C. Forlines, and F. Vernier. Beyond “social protocols”: Multi-user coordination policies for co-located groupware. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW '04)*, pages 262–265, New York, NY, USA, 2004. ACM Press.

- [194] C. Mühl, E. L. Broek, A.-M. Brouwer, F. Nijboer, N. Wouwe, and D. Heylen. Multi-modal affect induction for affective brain-computer interfaces. In S. D’Mello, A. Graesser, B. Schuller, and J.-C. Martin, editors, *Affective Computing and Intelligent Interaction*, volume 6974 of *LNCS*, pages 235–245. Springer Berlin Heidelberg, 2011.
- [195] M. A. Nacenta, D. Pinelle, D. Stuckel, and C. Gutwin. The effects of interaction technique on coordination in tabletop groupware. In *Proceedings of Graphics Interface (GI ’07)*, pages 191–198, New York, NY, USA, 2007. ACM Press.
- [196] C. Nass, I.-M. Jonsson, H. Harris, B. Reaves, J. Endo, S. Brave, and L. Takayama. Improving automotive safety by pairing driver emotion and car voice emotion. In *Extended Abstracts on Human Factors in Computing Systems (CHI EA ’05)*, pages 1973–1976, New York, NY, USA, 2005. ACM Press.
- [197] A. Neyem, C. Aracena, C. A. Collazos, and R. Alarcón. Designing emotional awareness devices: What one sees is what one feels. *Ingeniare Revista chilena de ingenieria*, 15(3), 2007.
- [198] R. Niño Zambrano and Y. Engelhardt. Diagrams for the masses: Raising public awareness – from neurath to gapminder and google earth. In G. Stapleton, J. Howse, and J. Lee, editors, *Diagrammatic Representation and Inference*, volume 5223 of *LNCS*, pages 282–292. Springer Berlin Heidelberg, 2008.
- [199] F. Nijboer, S. Carmien, E. Leon, F. Morin, R. Koene, and U. Hoffmann. Affective brain-computer interfaces: Psychophysiological markers of emotion in healthy persons and in persons with amyotrophic lateral sclerosis. In J. Cohn, A. Nijholt, and M. Pantic, editors, *Proceedings of the 3rd International Conference on Affective Computing and Intelligent Interaction (ACII ’09)*, pages 1–11, Los Alamitos, 2009. IEEE Computer Society Press.
- [200] A. Nijholt and D. Tan. Playing with your brain: Brain-computer interfaces and games. In *Proceedings of the International Conference on Advances in Computer Entertainment Technology (ACE ’07)*, pages 305–306, New York, NY, USA, 2007. ACM Press.
- [201] C. North. Toward measuring visualization insight. *IEEE Computer Graphics and Applications*, 26(3):6–9, 2006.

- [202] J. Ohene-Djan, A. Sammon, and R. Shipsey. Colour spectrum's of opinion: An information visualisation interface for representing degrees of emotion in real time. In *Proceedings of the International Conference on Information Visualization (IV '06)*, pages 80–88, Washington, DC, USA, 2006. IEEE Computer Society.
- [203] J. Ohene-Djan, J. Wright, and K. Combie-Smith. Emotional subtitles: A system and potential applications for deaf and hearing impaired people. In *Proceedings of the Conference and Workshop on Assistive Technologies for People with Vision and Hearing Impairments: Assistive Technology for All Ages (CVHI '07)*, 2007.
- [204] A. P. Oliveira and A. Cardoso. Towards affective-psychophysiological foundations for music production. In *Proceedings of the 2nd International Conference on Affective Computing and Intelligent Interaction (ACII '07)*, pages 511–522, Berlin, Heidelberg, 2007. Springer-Verlag.
- [205] T. Ormerod, J. Macgregor, and E. Chronicle. Dynamics and constraints in insight problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28:791–799, 2002.
- [206] J. Panksepp. A critical role for affective neuroscience in resolving what is basic about basic emotions. *Psychological Review*, 99(3):554–560, 1992.
- [207] B. Parker. *Aha!* Trafford Publishing, 2006.
- [208] J. Patten, H. Ishii, J. Hines, and G. Pangaro. Sensetable: a wireless object tracking platform for tangible user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*, pages 253–260, New York, NY, USA, 2001. ACM Press.
- [209] J. Patten, B. Recht, and H. Ishii. Audiopad: a tag-based interface for musical performance. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME '02)*, pages 1–6. National University of Singapore, 2002.

- [210] E. M. M. Peck, B. F. Yuksel, A. Ottley, R. J. Jacob, and R. Chang. Using fNIRS brain sensing to evaluate information visualization interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, pages 473–482, New York, NY, USA, 2013. ACM.
- [211] M. Perakakis and A. Potamianos. Affective evaluation of a mobile multimodal dialogue system using brain signals. In *Proceedings of the IEEE Spoken Language Technology Workshop (SLT '12)*, pages 43–48, 2012.
- [212] M. Perakakis and A. Potamianos. An affective evaluation tool using brain signals. In *Proceedings of the Companion Publication of the 2013 International Conference on Intelligent User Interfaces Companion (IUI '13 Companion)*, pages 105–106, New York, NY, USA, 2013. ACM.
- [213] P. Petrantonakis and L. Hadjileontiadis. Emotion recognition from brain signals using hybrid adaptive filtering and higher order crossings analysis. *IEEE Transactions on Affective Computing*, 1(2):81–97, July 2010.
- [214] R. Picard. *Affective Computing*. MIT Press, 1997.
- [215] C. Pindat, E. Pietriga, O. Chapuis, and C. Puech. JellyLens: Content-aware adaptive lenses. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology (UIST '12)*, pages 261–270, New York, NY, USA, 2012. ACM Press.
- [216] C. Plaisant. The challenge of information visualization evaluation. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '04)*, pages 109–116. ACM Press, 2004.
- [217] R. Plutchik. The nature of emotions. *American Scientist*, 89(4):344, 2001.
- [218] A. T. Pope, E. H. Bogart, and D. S. Bartolome. Biocybernetic system evaluates indices of operator engagement in automated task. *Biological Psychology*, 40(1–2):187–195, 1995.
- [219] S. Poslad. *Ubiquitous Computing: Smart Devices, Environments and Interactions*. Wiley, 1 edition, 2009.
- [220] V. Pravdich-Neminsky. Ein versuch der registrierung der elektrischen gehirnerscheinungen. *Zentralblatt für Physiologie*, 27:951–960, 1913.

- [221] W. Puff. *Towards a Collaborative Information Visualization System in a Multi-Display Environment*. PhD thesis, Institute for Computer Graphics and Vision, Graz University of Technology, 2009.
- [222] J. Qiu and Q. L. Zhang. “aha!” effects in a guessing chinese logogrph task: An event-related potential study. *Chinese Science Bulletin*, 53(3):384–391, 2008.
- [223] G. N. Ranky and S. Adamovich. Analysis of a commercial EEG device for the control of a robot arm. In *Proceedings of the 36th Annual IEEE Northeast Bioengineering Conference*, pages 1–2, 2010.
- [224] J. Reeder, J. Reeder, and R. Giegerich. Locomotif: From graphical motif description to RNA motif search in bioinformatics. *Bioinformatics*, 23:392–400, 2007.
- [225] J. Rekimoto, B. Ullmer, and H. Oba. DataTiles: A modular platform for mixed physical and graphical interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*, pages 269–276, New York, NY, USA, 2001. ACM Press.
- [226] L. Renambot, B. Jeong, R. Jagodic, A. Johnson, and J. Leigh. Collaborative visualization using high-resolution tiled displays. In *In ACM CHI '06 Workshop on Information Visualization, Interaction Techniques for Collaboration Across Multiple Displays*, 2006.
- [227] C. Reynolds and R. W. Picard. Evaluation of affective computing systems from a dimensional metaethical position. In *Proceedings of the 1st Augmented Cognition Conference (ACC '05)*, pages 22–27, 2005.
- [228] Y. Rogers and S. Lindley. Collaborating around vertical and horizontal large interactive displays: which way is best? *Interacting with Computers*, 16(6), 2004.
- [229] G. Ruffini, S. Dunnea, L. Fuentemillab, C. Graub, E. Farres, J. Marco-Pallares, P. Watts, and S. Silva. First human trials of a dry electrophysiology sensor using a carbon nanotube array interface. *Sensors and Actuators*, 144:275–279, 2008.
- [230] J. Russell. A circumplex model of affect. *Journal of Personality and Social Psychology*, 39:1161–1178, 1980.

- [231] J. Russell. Core affect and the psychological construction of emotion. *Psychological Review*, 110:145–172, 2003.
- [232] J. Russell and L. Feldman Barrett. Core affect. In D. Sander and K. Scherer, editors, *The Oxford companion to emotion and the affective sciences*. Oxford University Press, 2009.
- [233] J. A. Russell and L. F. Barrett. Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant. *Journal of Personality and Social Psychology*, 76(5):805–819, 1999.
- [234] K. Ryberg. *Levande färger*. ICA Bokförlag, Västerås, Sweden, 1999.
- [235] T. Saari, K. Kallinen, M. Salminen, N. Ravaja, and M. Rapino. A mobile and desktop application for enhancing group awareness in knowledge work teams. In J. Jacko, editor, *Human-Computer Interaction. Ambient, Ubiquitous and Intelligent Interaction*, volume 5612 of *LNCS*, pages 95–104. Springer Berlin Heidelberg, 2009.
- [236] T. Saari, K. Kallinen, M. Salminen, N. Ravaja, and K. Yanev. A mobile system and application for facilitating emotional awareness in knowledge work teams. In *Proceedings of the 41st Annual Hawaii International Conference on System Sciences (HICSS '08)*, pages 44–53, 2008.
- [237] N. Sakai. The role of sentence closing as an emotional marker: A case of japanese mobile phone e-mail. *Discourse, Context & Media*, 2(3):149–155, 2013.
- [238] P. Salovey and J. D. Mayer. Emotional intelligence. *Imagination, Cognition, and Personality*, 9:185–211, 1990.
- [239] P. Saraiya, C. North, and K. Duca. An insight-based methodology for evaluating bioinformatics visualizations. *IEEE Transactions on Visualizations and Computer Graphics*, 11(4):443–456, 2005.
- [240] M. Scheerer. Problem-solving. *Scientific American*, 208:118–128, 1963.
- [241] K. Scherer. Toward a dynamic theory of emotion: The component process model of affective states. *Geneva Studies in Emotion and Communication*, 1:1–98, 1987.

- [242] K. Scherer. Psychological models of emotion. *The Neuropsychology of Emotion*, pages 137–162, 2000.
- [243] K. Scherer. What are emotions? and how can they be measured? *Social Science Information*, 44(4):695–729, December 2005.
- [244] R. Scherer, G. Muller, C. Neuper, B. Graimann, and G. Pfurtscheller. An asynchronously controlled EEG-based virtual keyboard: improvement of the spelling rate. *IEEE Transactions on Biomedical Engineering*, 51(6):979–984, 2004.
- [245] B. Schiettecatte and J. Vanderdonckt. AudioCubes: a distributed cube tangible interface based on interaction range for sound design. In *Tangible and Embedded Interaction*, pages 3–10. ACM Press, 2008.
- [246] D. Schmalstieg, L. M. Encarnação, and Z. Szalavári. Using transparent props for interaction with the virtual table. In *Proceedings of the 1999 Symposium on Interactive 3D Graphics (I3D '99)*, pages 147–153, New York, NY, USA, 1999. ACM Press.
- [247] F. Schmidt. *Collaborative Motif Search in Large Browser History Visualizations through Combinable Tangible Rings*. B.sc. thesis, University of Kaiserslautern, Germany, 2013.
- [248] H. Schmidt, U. Hinrichs, A. Dunning, and S. Carpendale. Memory [En]Code: Building a collective memory within a tabletop installation. In *Proceedings of the 3rd Eurographics Conference on Computational Aesthetics in Graphics, Visualization and Imaging (Computational Aesthetics '07)*, pages 135–142, Aire-la-Ville, Switzerland, 2007. Eurographics Association.
- [249] S. D. Scott, M. S. T. Carpendale, and S. Habelski. Storage bins: Mobile storage for collaborative tabletop displays. *IEEE Computer Graphics and Applications*, 25(4):58–65, July 2005.
- [250] S. D. Scott, M. S. T. Carpendale, and K. M. Inkpen. Territoriality in collaborative tabletop workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW '04)*, pages 294–303, New York, NY, USA, 2004. ACM Press.

- [251] S. D. Scott, K. D. Grant, and R. L. Mandryk. System guidelines for co-located, collaborative work on a tabletop display. In *Proceedings of the 8th Conference on European Conference on Computer Supported Cooperative Work (ECSCW '03)*, pages 159–178, Norwell, MA, USA, 2003. Kluwer Academic Publishers.
- [252] P. Sengers, K. Boehner, S. Warner, and T. Jenkins. Evaluating affector: Co-interpreting what works. In *CHI 2005 Workshop on Innovative Approaches to Evaluating Affective Systems*, 2005.
- [253] O. Shaer and E. Hornecker. Tangible user interfaces: Past, present, and future directions. *Foundations and Trends in Human-Computer Interaction*, 3(1):1–137, January 2010.
- [254] F. Sharbrough, G. Chatrian, R. Lesser, H. Lüders, M. Nuwer, and T. Picton. American electroencephalographic society guidelines for standard electrode position nomenclature. *Journal of Clinical Neurophysiology*, 8:200–202, 1991.
- [255] S.-T. Shen, S. Prior, and K.-M. Chen. A solution to revisitation using organic bookmark management. In A. Marcus, editor, *Design, User Experience, and Usability. Theory, Methods, Tools and Practice*, volume 6770 of *LNCS*, pages 46–52. Springer Berlin Heidelberg, 2011.
- [256] B. Sheth, S. Sandkühler, and J. Bhattacharya. Posterior beta and anterior gamma oscillations predict cognitive insight. *Journal of Cognitive Neuroscience*, 21(7):1269–1279, 2009.
- [257] M. N. Shiota and J. W. Kalat. *Emotion*. Cengage Learning, 2 edition, 2011.
- [258] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the 1996 IEEE Symposium on Visual Languages (VL '96)*, pages 336–343, Washington, DC, USA, 1996. IEEE Computer Society.
- [259] B. Shneiderman. Computer Science: Science 2.0. *Science*, 319:1349–1350, 2008.
- [260] Y. B. Shrinivasan and J. J. van Wijk. Supporting exploration awareness in information visualization. *IEEE Computer Graphics and Applications*, 29(5):34–43, September 2009.

- [261] M. Shugrina, M. Betke, and J. Collomosse. Empathic painting: Interactive stylization through observed emotional state. In *Proceedings of the 4th international Symposium on Non-photorealistic Animation and Rendering (NPAR '06)*, pages 87–96. ACM Press, 2006.
- [262] O. Sourina and Y. Liu. EEG-enabled affective human-computer interfaces. *Universal Access in Human-Computer Interaction. Design and Development Methods for Universal Access*, 8513:536–547, 2014.
- [263] M. Spindler, S. Stellmach, and R. Dachsel. PaperLens: Advanced magic lens interaction above the tabletop. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '09)*, pages 69–76, New York, NY, USA, 2009. ACM Press.
- [264] M. Spindler, C. Tominski, H. Schumann, and R. Dachsel. Tangible views for information visualization. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '10)*, pages 157–166, New York, NY, USA, 2010. ACM Press.
- [265] A. Spink, M. Park, B. J. Jansen, and J. Pedersen. Multitasking during web search sessions. *Information Processing & Management*, 42(1):264–275, 2006.
- [266] A. Ståhl, K. Höök, M. Svensson, A. S. Taylor, and M. Combetto. Experiencing the affective diary. *Personal and Ubiquitous Computing*, 13(5):365–378, 2009.
- [267] A. Ståhl, P. Sundström, and K. Höök. A foundation for emotional expressivity. In *Proceedings of the 2005 Conference on Designing for User eXperience (DUX '05)*, page Article 33, New York, NY, USA, 2005. AIGA: American Institute of Graphic Arts.
- [268] J. T. Stasko and Q. A. Zhao. Visualizing the execution of threads-based parallel programs. Technical Report GIT-GVU-95-01, Georgia Institute of Technology, 1995.
- [269] E. Steiner, A. MacEachren, and D. Guo. Developing lightweight, data-driven exploratory geo-visualization tools for the web. In D. E. Richardson and P. van Oosterom, editors, *Advances in Spatial Data Handling*, pages 487–500. Springer Berlin Heidelberg, 2002.

- [270] D. Stone, C. Jarrett, M. Woodroffe, and S. Minocha. *User Interface Design and Evaluation*. Morgan Kaufmann/Elsevier, 2005.
- [271] M. C. Stone, K. Fishkin, and E. A. Bier. The movable filter as a user interface tool. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '94)*, pages 306–312, New York, NY, USA, 1994. ACM Press.
- [272] K. Stytsenko, E. Jablonskis, and C. Prahm. Evaluation of consumer EEG device Emotiv EPOC. In *Proceedings of the MEi CogSci Conference, Poster*, Ljubljana, Slovenia, 2011.
- [273] K. Suchner. *EmotionWidgets: Emotionen in GUI Controls*. B.sc. thesis, University of Kaiserslautern, Germany, 2012.
- [274] P. Sundström, T. Jaensson, K. Höök, and A. Pommeranz. Probing the potential of non-verbal group communication. In *Proceedings of the ACM 2009 International Conference on Supporting Group Work (GROUP '09)*, pages 351–360, New York, NY, USA, 2009. ACM Press.
- [275] B. Swartz and E. Goldensohn. Timeline of the history of EEG and associated fields. *Electroencephalography and Clinical Neurophysiology*, 106(2):173–176, 1998.
- [276] A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale. Collaborative coupling over tabletop displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*, pages 1181–1190, New York, NY, USA, 2006. ACM Press.
- [277] G. S. Taylor and C. Schmidt. Empirical evaluation of the Emotiv EPOC BCI headset for the detection of mental actions. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 56(1):193–197, 2012.
- [278] S. Thelen, J. Meyer, A. Middel, P. Olech, A. Ebert, and H. Hagen. Tag-based interaction with large high-resolution displays. In *Proceedings of the 4th IASTED International Conference on Human-Computer Interaction (IASTED-HCI)*, 2009.
- [279] J. Thomas and K. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. National Visualization and Analytics Center, 2005.

- [280] M. Thuring and S. Mahlke. Usability, aesthetics and emotions in human-technology interaction. *International Journal of Psychology*, 42(4):253–264, 2007.
- [281] M. Tobiasz, P. Isenberg, and S. Carpendale. Lark: Coordinating co-located collaboration with information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1065–1072, November 2009.
- [282] C. Tominski, S. Gladisch, U. Kister, R. Dachsel, and H. Schumann. A survey on interactive lenses in visualization. In *EuroVis State-of-the-Art Reports*, pages 43–62. Eurographics Association, 2014.
- [283] M. Tory and T. Möller. Human factors in visualization research. *IEEE Transactions on Visualization and Computer Graphics*, 10(1):72–84, January 2004.
- [284] A. Treisman. Features and objects in visual processing. *Scientific American*, 255(5):114–125, 1986.
- [285] I. Truderung. *Interaktive Visualisierung erweiterter Browserchroniken*. B.sc. thesis, University of Kaiserslautern, Germany, 2011.
- [286] B. Ullmer and H. Ishii. The metaDESK: models and prototypes for tangible user interfaces. In *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology (UIST '97)*, pages 223–232, New York, NY, USA, 1997. ACM Press.
- [287] B. Ullmer, H. Ishii, and R. Jacob. Token+constraint systems for tangible interaction with digital information. *ACM Transactions on Computer-Human Interaction*, 12(1):81–118, 2005.
- [288] J. Underkoffler and H. Ishii. Illuminating light: an optical design tool with a luminous-tangible interface. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '98)*, pages 542–549, New York, NY, USA, 1998. ACM Press/Addison-Wesley Publishing Co.
- [289] J. Underkoffler and H. Ishii. Urp: A luminous-tangible workbench for urban planning and design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*, pages 386–393, New York, NY, USA, 1999. ACM Press.

- [290] M. Van De Velde, G. Van Erp, and P. Cluitmans. Detection of muscle artefact in the normal human awake EEG. *Electroencephalography and Clinical Neurophysiology*, 107(2):149–158, 1998.
- [291] G. C. van der Veer, O. Kulyk, and D. Vyas. Mirrors of the World – supporting situational awareness with computer screens. In D. Gorgan and A. M. Guran, editors, *Interactiune Om-Calculator (RoCHI '09)*, Bucharest, Romania, September 2009. Matrix Rom.
- [292] M. Van Kleek, B. Moore, C. Xu, and D. R. Karger. Eyebrowse: Real-time web activity sharing and visualization. In *Extended Abstracts on Human Factors in Computing Systems (CHI EA '10)*, pages 3643–3648, New York, NY, USA, 2010. ACM Press.
- [293] F. B. Viégas, M. Wattenberg, F. van Ham, J. Kriss, and M. McKeon. ManyEyes: A site for visualization at internet scale. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1121–1128, 2007.
- [294] M. Viermetz, C. Stolz, V. Gedov, and M. Skubacz. Relevance and impact of tabbed browsing behavior on web usage mining. In *Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence (WI '06)*, pages 262–269, 2006.
- [295] S. Voelker, K. Nakajima, C. Thoresen, Y. Itoh, K. I. Øvergård, and J. Borchers. PUCs: Detecting transparent, passive untouched capacitive widgets on unmodified multi-touch displays. In *Proceedings of the Adjunct Publication of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13 Adjunct)*, pages 1–2, New York, NY, USA, 2013. ACM Press.
- [296] T. von Landesberger, R. Rehner, M. Görner, and T. Schreck. A system for interactive visual analysis of large graphs using motifs in graph editing and aggregation. In *Proceedings of the Vision Modeling Visualization Workshop (VMV '09)*, pages 331–340, 2009.
- [297] U. von Zadow, F. Daiber, J. Schöning, and A. Krüger. GlobalData: Multi-user interaction with geographic information systems on interactive surfaces. In M. Haller and D. Wigdor, editors, *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '10)*. ACM Press, 2010.

- [298] M. Waldner, J. Hauber, J. Zauner, M. Haller, and M. Billinghamurst. Tangible tiles: design and evaluation of a tangible user interface in a collaborative tabletop setup. In *Proceedings of the 18th Australia Conference on Computer-Human Interaction: Design: Activities, Artifacts and Environments (OZCHI '06)*, pages 151–158, New York, NY, USA, 2006. ACM Press.
- [299] M. Waldner, A. Lex, M. Streit, and D. Schmalstieg. Design considerations for collaborative information workspaces in multi-display environments. In *Proceedings of Workshop on Collaborative Visualization on Interactive Surfaces (CoVIS '09), in conjunction with VisWeek*, 2009.
- [300] J. B. Walther and K. P. D’Addario. The impacts of emoticons on message interpretation in computer-mediated communication. *Social Science Computer Review*, 19(3):324–347, 2001.
- [301] B. Wang, N. Zhang, J. Hu, and Z. Shen. Visualizing web browsing history with barcode chart. In *Proceedings of the Adjunct Publication of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13 Adjunct)*, pages 99–100, New York, NY, USA, 2013. ACM Press.
- [302] X. Wang, W. Dou, W. Ribarsky, D. Skau, and M. X. Zhou. LeadLine: Interactive visual analysis of text data through event identification and exploration. In *Proceedings of the 2012 IEEE Conference on Visual Analytics Science and Technology (VAST '12)*, pages 93–102, Washington, DC, USA, 2012. IEEE Computer Society.
- [303] F. Wanner, A. Weiler, and T. Schreck. Topic Tracker: Shape-based visualization for trend and sentiment tracking in twitter. In *Proceedings of the 2nd IEEE Workshop on Interactive Visual Text Analytics as part of the IEEE VisWeek 2012*, 2012.
- [304] C. Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers, 2 edition, 2004.
- [305] J. Waser, R. Fuchs, H. Ribicic, B. Schindler, G. Bloschl, and E. Gröller. World lines. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1458–1467, 2010.

- [306] S. Waterson, J. I. Hong, T. Sohn, J. Heer, T. Matthews, and J. Landay. What did they do? understanding clickstreams with the WebQuilt visualization system. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '02)*, pages 94–102, 2002.
- [307] H. Weinreich, H. Obendorf, E. Herder, and M. Mayer. Off the beaten tracks: exploring three aspects of web navigation. In *Proceedings of the 15th International Conference on World Wide Web (WWW '06)*, pages 133–142, 2006.
- [308] M. Weiss, J. Wagner, R. Jennings, Y. Jansen, R. Khoshabeh, J. D. Hollan, and J. Borchers. SLAPbook: tangible widgets on multi-touch tables in groupware environments. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction (TEI '09)*, pages 297–300, New York, NY, USA, 2009. ACM Press.
- [309] P. Wellner. Interacting with paper on the digitaldesk. *Communications of ACM*, 36(7):87–96, July 1993.
- [310] D. Wigdor, C. Shen, C. Forlines, and R. Balakrishnan. Perception of elementary graphical elements in tabletop and multi-surface environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*, pages 473–482, New York, NY, USA, 2007. ACM.
- [311] D. Wigdor, S. Williams, M. Cronin, R. Levy, K. White, M. Mazeev, and H. Benko. Ripples: Utilizing per-contact visualizations to improve user interaction with touch displays. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology (UIST '09)*, pages 3–12, New York, NY, USA, 2009. ACM Press.
- [312] D. Wijayasekara and M. Manic. Human machine interaction via brain activity monitoring. In *Proceedings of the 6th International Conference on Human System Interaction (HSI '13)*, pages 103–109, June 2013.
- [313] W. Willett, J. Heer, and M. Agrawala. Scented Widgets: Improving navigation cues with embedded visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1129–1136, 2007.
- [314] T. Wills, S. Soraci, R. Chechile, and H. Taylor. “aha” effects in the generation of pictures. *Memory & Cognition* 28 (6), 28(6):939–948, 2000.

- [315] J. Wood, H. Wright, and K. Brodlie. Collaborative visualization. In *Proceedings of the IEEE Symposium on Visualization (Vis '97)*, pages 253–259, 1997.
- [316] W. Yu and T. Ingalls. Trails – an interactive web history visualization and tagging tool. In A. Marcus, editor, *Design, User Experience, and Usability. Theory, Methods, Tools and Practice*, volume 6770 of *LNCS*, pages 77–86. Springer Berlin Heidelberg, 2011.
- [317] T. O. Zander, C. Kothe, S. Jatzev, and M. Gaertner. Enhancing human-computer interaction with input from active and passive brain-computer interfaces. In *Brain-Computer Interfaces*, Human-Computer Interaction Series, pages 181–199. Springer London, 2010.
- [318] S. Zhang, Q. Huang, S. Jiang, W. Gao, and Q. Tian. Affective visualization and retrieval for music video. *Transactions on Multimedia*, 12(6):510–522, October 2010.
- [319] N. N. Zhu and S. Q. Sun. Visual thinking in data visualization. In W. Deng and Q. Luo, editors, *Applied Mechanics and Materials*, volume 236–237, pages 597–602. Trans Tech Publications, Switzerland, 2012.
- [320] G. Zufferey, P. Jermann, A. Lucchi, and P. Dillenbourg. TinkerSheets: Using paper forms to control and visualize tangible simulations. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction (TEI '09)*, pages 377–384, New York, NY, USA, 2008. ACM Press.

Linnaeus University Dissertations

Below please find a list of recent publications in the series Linnaeus University Dissertations. For a full list and more information: Lnu.se

206. Ke Yang 2015, *Regulation of artemisinin biosynthesis in Artemisia annua L* (biomedicin/biomedical science) ISBN: 978-91-87925-43-6.

207. Henrik Nilsson 2015, *Kultur och utbildning – en tolkning av två grundskolors mångkulturella kontexter* (pedagogik/education) ISBN: 978-91-87925-44-3.

208. Ida Ohlsson Al Fakir 2015, *Nya rum för socialt medborgarskap. Om vetenskap och politik i "Zigenarundersökningen" – en socialmedicinsk studie av svenska romer 1962-1965* (historia/history) ISBN: 978-91-87925-46-7.

209. Antonina Danylenko 2015, *Decision Algebra: A General Approach to Learning and Using Classifiers* (datavetenskap/computer science) ISBN: 978-91-87925-47-4.

210. René Nordstierna 2015, *Osäkra möten. Perspektiv på metodbildning för det konstnärliga högskoleutbildningsfältet* (pedagogik/education) ISBN: 978-91-87925-48-1.

211. Reza Mahani 2015, *Magnetic solotronics near the surface of a semiconductor and a topological insulator* (fysik/physics) ISBN: 978-91-87925-49-8.

212. Niklas Andersson 2015, *När vårdande och lärande sammanfaller. Patienters, studenter och handledares erfarenheter av möten på en utbildningsvårdavdelning inom psykiatrisk vård* (vårdvetenskap/nursing) ISBN: 978-91-87925-50-4.

213. Luca Fumarco 2015, *Essays on discrimination in the marketplace* (nationalekonomi/economics) ISBN: 978-91-87925-51-1.

214. Kerstin Hemström 2015, *Perceptions of intensive forestry and multi-storey wood frames among Swedish actors* (byggteknik/civil engineering) ISBN: 978-91-87925-52-8.

215. Maria Alm 2015, *När kunskap ska styra – om organisatoriska och professionella villkor för kunskapsstyrning inom missbruksvården* (pedagogik/education) ISBN: 978-91-87925-54-2.

216. Anna Sävneby 2015, *Reverse genetic studies of Enterovirus replication* (biomedicinsk vetenskap/biomedical science) ISBN: 978-91-87925-55-9.

217. Michelle Wille 2015, *Viruses on the wing: evolution and dynamics of influenza A virus in the Mallard reservoir* (ekologi/ecology) ISBN: 978-91-87925-56-6.

218. Hiba Nassar 2015, *The Hodrick-Prescott Filter: Functional aspects and statistical estimation* (matematik/mathematics) ISBN: 978-91-87925-57-3.
219. Andreas Ivarsson 2015, *Psychology of sport injury: Prediction, prevention and rehabilitation in Swedish team sport athletes* (psykologi/psychology) ISBN: 978-91-87925-58-0.
220. Lina Palmér 2015, *Amning och Existens. Moderskap, sårbarhet och ömsesidigt beroende vid inledande amning* (vårdvetenskap/nursing) ISBN: 978-91-87925-59-7
221. Ann-Charlotte Norman 2015, *Towards the creation of learning improvement practices: Studies of pedagogical conditions when change is negotiated in contemporary healthcare practices* (pedagogik/ education) ISBN: 978-91-87925-61-0.
222. Petter Tibblin, 2015, *Migratory behaviour and adaptive divergence in life-history traits of pike (*Esox lucius*)* (ekologi/ecology) ISBN: 978-91-87925-62-7.
223. Emma Neuman, 2015, *Essays on Segregation, Gender Economics, and Self-employment* (nationalekonomi/economics) ISBN: 978-91-87925-63-4.
224. Daniel Cernea, 2015, *User-Centered Collaborative Visualization* (datavetenskap/computer science) ISBN: 978-91-87925-65-8.