Enhancing Collaboration Through Role Specific Information Sharing

by

Juliano Franz

Submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy

at

Dalhousie University
Halifax, Nova Scotia
January 2023

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

Information sharing, either directly or trough summaries and proxies, is one of the pillars of collaborative systems in human-computer interaction. Representing and sharing role-specific information in collocated systems is one of the many challenges in groupware systems. In my thesis, I explore methods and their impacts on sharing role-based information in applications where users do not have uniform access to information either because they are using different technologies to access data or because there is no support for the flow of information. I present tools spanning both "What You See Is What I See" and "What You See Is Not What I See" approaches displaying either raw or summarized information to users. I present three studies leading towards an approach to support road cyclists to perform better as a unit by sharing their exertion information among all group members while leveraging technology ubiquitous to them. I start my work with a study exploring how to enable shared augmented reality experiences in museums when only one of the group members has access to an augmented reality headset. Either because there are not enough devices available or someone might feel uncomfortable using one in public. I propose two approaches for AR experiences (Over-the-Shoulder AR and Semantic Linking), a complementary technique (Indicator Rings), evaluate them in a long-term in-the-wild study, and discuss their impacts on museum scenarios and other applications. Later I present a design research study exploring the work done by road cyclists while training together, following a Contextual Design methodology. I expose their strategies and challenges with group coordination and communication and uncover a lack of support for sharing individual metrics such as effort in standard tools. Finally, I present my last study comparing the impacts on performance and self-reported metrics of two methods for road cyclists to share their exertion levels with group members while exercising: Paceguide and RPE View. I then discuss the broader implications of my work for the HCI community while proposing future research venues in Sport-HCI and other co-located work domains.

LIST OF ABBREVIATIONS USED

%FTP: Percentage of Functional Threshold Power

BLE: Bluetooth Low Energy

BtF: Back to front rotation

CD: Contextual Design

FtB: Front to back rotation

FTP: Functional Threshold Power

HR: Heart Rate

HWD: Head Worn Display

IR: Indicator Rings

LTD: Laurens Ten Dam

MMA: Maritime Museum of the Atlantic

NP: Normalized Power

OtS: Over-the-Shoulder AR

RPE: Ratings of Perceived Exertion

SL: Semantic Linking

SUS: System Usability Scale

TTT: Team Time Trial

W/kg: Watts per kilo

WIMP: Windows Icons Mouse and Pointers

WYSINWIS: What You See is Not What I See

WYSIWIS: What You See is What I See

ACKOWLEDGEMENTS

First, I would like to offer my gratitude to my supervisor, Dr. Derek Reilly for his guidance and kindship. His support was paramount in inspiring me and helping me to complete my research work.

I want to thank my committee members Dr. Fernando Paulovitch and Dr. Paul Ralph for their time and expertise. Special thanks to Professor Carman Neustaedter for accepting to be my external examiner.

I would also like to thank Dr. Joseph Malloch for his hours of insightful discussions throughout my program and for his questions during the thesis defence. Thank you, GEM Lab members, for keeping me sane throughout all these years. Thank you to my family back home for their unyielding support and understanding of my decision to move far away to pursue my studies. A long overdue thank you to Dr. Luciana Nedel for guiding me through my undergraduate research days and opening my mind to the research world!

Lastly, I would like to express my love and gratitude to my wife Amanda for her patience and support throughout the Ph.D. program and to my baby daughter Clarissa for making sure I was awake early to work on my thesis (albeit not rested enough).

CHAPTER 1 INTRODUCTION

Many modern digital devices are moving away from the traditional WIMP desktop interaction models (Roberts et al., 2014) and moving into more mobile forms of presentation and interaction, such as smartphones, tablets, small displays, and augmented-reality headsets. As these devices become ubiquitous or are introduced in domain-specific tasks, it is only natural to expect them to be integral to face-to-face or remote collaborative systems.

Devices such as augmented reality headsets or hand-held displays enable the possibility of interacting with 2D and 3D data using gestures, voice commands and body movements. Paired with registration algorithms, virtual objects can be placed within spaces in the real world, creating integrated experiences that leverage elements of physicality and virtuality. When combined with wireless networks, digital experiences can be shared across multiple devices, including different display technologies such as projectors, large wall displays and head-worn displays.

This new mixed ecology creates an interesting research space on how to support collaboration when collocated members have access to different display technologies, multiple roles, and different access restrictions to information. As information access is now not uniform, there is a need to create grounding references (such as a shared visualization space) to support conversation and collaboration.

When access to visualization devices is not uniform, enabling sharing of role-specific information can enhance engagement as it has the potential to close the gap between two or more people trying to coordinate work. This can be done through direct symmetrical sharing of a specific view using (relaxed)WYSIWIS approaches or through summaries or proxies for information with asymmetrical WYSINWIS designs. The advantages and disadvantages of sharing role-specific information to promote engagement in collaborative systems using WYSIWIS vs WYSINWIS and literal/raw vs. summarized data warrant exploration.

Throughout my thesis, I explore the role of collaborative digital systems in bringing people together toward a shared understanding of their situation. My work revolves around three studies exploring collocated collaboration through symmetrical and asymmetrical sharing when users with different roles are exploring data together. In the following chapters, I follow the answer ramifications of the following question: What are the impacts of sharing of role specific data views on communication and coordination of collocated collaborative spaces?

In the first study, presented in Chapter 3, I explore how much sharing is needed to bring museum visitors towards the understanding of a mixed reality exhibit while only one person has access to an augmented reality device. This work serves as one of the foundation points for understanding sharing and communication between people working together but with different roles and various levels of access to information, the common thread throughout my Ph.D. Chapter 5 presents the exploration and design work with my target population, road cyclists. I expose the contextual design work that lays the ground for an in-depth understanding of the cyclists' needs, internal and external processes, and attrition points while they work together with other riders when exercising. Lastly, in Chapter 6, I tackle the lack of effort sharing among riders in a group as of the issues identified in

Chapter 5; I validate the impacts of a set of exertion sharing tools on the rider's behaviour and performance metrics with and without the tools.

1.1 Contribution to the Field

Challenges within the CSCW domain are vast and include communication methods, data sensing and sharing, multi-role support, and group coordination. My thesis work, presented in the following chapters, aims to design and evaluate tools to promote awareness and coordination in collocated groups, enabling users to see and share role-specific information while maintaining some form of common reference available to everyone to facilitate engagement. The domain-specific applications used as case studies in my thesis contribute with their own technical and dsocial challenges. For example, the need for polished engaging experiences supporting multiple users without technical knowledge or the stress and physical demand of the sports activity and the environmental factors such as adverse weather conditions.

The contributions I make on my thesis are as follows:

- A set of three novel sharing techniques for Mixed Reality applications to support collaboration and engagement with content between two or more people when only one person has access to the virtual content.
- A controlled longitudinal field study exploring the effectiveness of creating a joint exploration of a mixed reality exhibit using the sharing techniques mentioned above.
- An in-depth exploration of the collaborative group work done by road cyclists riding together. I focused on exposing their strategies to deal with their

environment, work tasks while actively exercising, strategies to deal with weather and hazards, and the challenges managing exertion of all riders while pacing the group.

- A collection of exertion sharing tools tailored for bicycle computers using the learned design outcomes from the aforementioned exploration work.
- An in-the-wild user evaluation of two exertion sharing visualizations comparing the performance of riders exposed to both techniques showing an increased power output and better awareness of their peers when compared to their traditional baseline.

CHAPTER 2 BACKGROUND AND PRELIMINARY WORK

In this chapter I briefly present higher-level overview of some of the grounding work of my thesis and then my preliminary work during my Ph.D. that led the framing of my thesis. Not all of this work was aligned with the main topic of my thesis, but it served as a formative process narrowing my research focus. I introduce a more in-depth, domain specific, analysis of the supporting research for my work on Chapters 3 and 4 paired with their respective user studies.

Information sharing to promote awareness is paramount in supporting collaboration between individuals and groups (Gutwin & Greenberg, 1996). How to capture and share information is a challenging topic as CSCW systems are required to support different interaction and presentation modalities depending on their application (Johansen, 1988).

What You See is What I See (WYSIWIS) approaches to information sharing enable users to see the same view independently of their roles or location. Strict WISIWYS, however is "inflexible" and does not differentiate between the needs of the user and the group (Stefik et al., 1987). In the other end of the spectrum, What You See is Not What I See (WISINWYS) interfaces (Zhu, 2004) focus on having tailored to the needs or requirements of specific individuals but may fail to provide some form of common grounding to promote engagement or conversation between collaborating members. *Relaxed* WYSIWIS systems allow for sharing of a common view like strict WYSIWIS but also support specific needs of users and groups enabling, for example, role specific visualizations (Nilsson et al., 2009) or asynchronous interactions.

Supporting multiple roles and their needs (individually and shared among the group) is also key in enabling collaborative experiences (Guzdial et al., 2000). When users with different roles collaborate, tools such as Boundary Objects[1] can facilitate communication between their members when required (Star & Griesemer, 1989). Both Road Cycling and Museum Experiences can be social activities (Hood, 1983; Warner et al., 2012); supporting their group work (either exploring an exhibit space or coordinating the practice of sport) can increase engagement and enjoyment.

2.1 Preliminary work leading to my thesis

I began my Ph.D. designing visualizations that synthesized time-sensitive data into a simple visual that allowed decision-makers to determine appropriate courses of action. This work was in the context of space satellite monitoring, wherein operators need to remain aware of trajectory data of a large number of satellites (both under their control and otherwise), and to detect anomalous circumstances that may require them to make timely decisions to alter orbits or speeds in order to avoid collisions.

This visualization design work directly informed my subsequent visualization work in naval information space and road cycling: in both cases, decisions are made that require the time-sensitive synthesis and presentation of heterogeneous data.

While this project did not proceed to the point where these visualizations could be formally evaluated, I did participate in feedback sessions with two operators. The operators were quite blunt in their appraisal of several of the proposed data visualization approaches:

[1] Standardized digital or physical artifacts used to facilitate the exchange of heterogenous information between parties from different domain knowledges.

specifically, visual monitoring was highly unlikely due to their work constraints; they prioritized notification mechanisms for that reason and were not interested in visual data synthesis unless it provided clear benefits in the course of specific decision-making scenarios. One of the key takeaways for me from these sessions was how critically important it was to build a deep understanding of work practice and the impacts interfaces and visual representations have on that work practice, including one's ability to make informed decisions. This shaped my research approach going forward: in particular, I employed a thorough Contextual Design process for my work in the road cycling domain.

After the space satellite visualization work, I led a research project with others in our lab exploring the use of compensatory techniques for perspective warping when close to very large 2D displays. My perspective compensation technique (Figure 1) maintains objects' relative size and shape as they move away from the viewer, thus mitigating the natural size distortion of far objects.

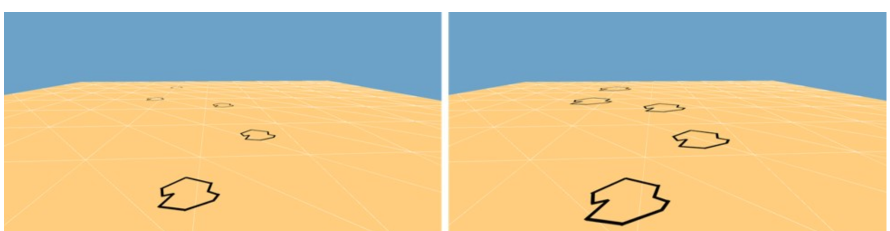

*Figure 1 -The effect of perspective compensation as perceived from an observer close to the surface.*Left: perspective compensation is off. Right: perspective compensation is on. '

The research combined controlled experiments with an in-the-wild study during a public art festival. This work was published as a paper (J. Franz et al., 2020) and presented as a demo (J. M. Franz et al., 2018). While the primary objectives of that research project are not directly related to my thesis, the work still impacted my thinking in significant ways. First, an interesting aspect of the work was that it allowed us to observe whether the visual compensation techniques – which involved resizing screen content based on the position and orientation of a single user – impacted the experience of bystanders watching what that single user was doing. Observations indicated that the crowd was not impacted in any obvious way when this compensation was active. While inconclusive, this finding encouraged me to further explore how data tailored for a single person can be shared with others to benefit understanding and collaborative work while still providing visual representations that were tailored to an individual's perspective, their role, and/or to the capacities and limitations of their digital devices. Second, several challenges and benefits of in-the-wild evaluations were made apparent to me in this study: these include challenges faced when running evaluations within time-bounded events, the varied influence of a crowd on participant engagement, and the benefits of integrating the study into the broader situation or context that it is a part of. This helped me to design an effective in-the-wild comparative evaluation in the museum context and to design an experimental simulation in the road cycling context with a high level of ecological validity.

In parallel to the work on perspective compensation, I started exploring how to support the decision-making process of multiple actors that either had different access levels to the same information due to different roles or were accessing information from different types of devices (e.g., AR headset vs tablet). Furthermore, I was also interested in exploring interaction challenges in Mixed Reality environments, such as how to select elements, support for natural visualization gestures (overview, zoom, filter, history) and touch feedback.

This work as part of the Mission-relevant Information Management for Integrated Response (MIMIR) project with the goal of supporting maritime domain information visualization from multiple perspectives to support the decision-making process while at sea. As an ideation tool, I created a prototype to provide multiple people access to Mixed Reality, which combined AR devices, tablets, and a projection surface to show geo-spatial information common to all actors. Figure 2 shows the early work exploring vessel communication data from navy vessels at sea. Two data specialists explore the information with different access levels while the shared elements, such as the map projection and data flow arches, create enough common grounding. This early work was published as a workshop paper (J. Franz, Malloch, Nedel, et al., 2017) and a short paper (Franz, Malloch, & Reilly, 2017).

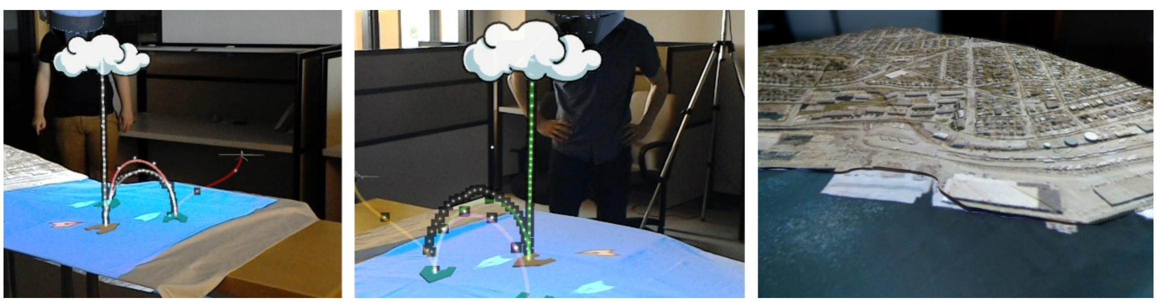

Figure 2 - Images on the left and center depict two users exploring communication flow between a group of vessels, drones and cloud data. Each user sees the data categorized by his role; however, some elements are constant, providing virtual spatial cues for face-to-face collaboration. The image on the right shows a detail of our static tangible tabletop display---the land contour is constructed from CNC-milled wood; aerial photographs or other imagery is added using projection mapping.

CHAPTER 3 A COMPARATIVE EVALUATION OF TECHNIQUES FOR SHARING AR EXPERIENCES IN MUSEUMS

I started my work on supporting collaboration between people with heterogeneous access to information inspired by the challenges in data visualization and manipulation faced by users in the Maritime Domain and the previous AR work developed by Schemalsteig et al (1996) and Vlahakis et al (2001). The focus was on interaction beyond the desktop as data collection has become increasingly more intense, and datasets are complex and connected, requiring novel interaction and visualization techniques (Chandler et al., 2015; Reda et al., 2013; Roberts et al., 2014).

The study presented in this chapter sets the common thread on my work, that is, supporting collaborative experiences supporting role specific interactions and visualizations. To fully explore this theme, one must focus on creating domain specific applications where I could get practical insights from observing and evaluating people performing real, specific tasks.

I refined prototype designed in the MIMIR (the naval work referenced in Section 2.1) project to explore effective means to share information with two people with heterogeneous access (AR vs no-AR) while promoting engagement between them in a museum setting. The museum work was a two-year collaboration between our research lab and the Narratives in Space and Time group[2], a group of artists, architects, and social scientists exploring the use of locative media and its role in raising community awareness of our shared urban landscape and the decisions that shape it. Together, we created the

[2] http://www.narrativesinspaceandtime.ca/

Psychogeographer's Table as part of the Halifax Explosion 100th Anniversary exhibition in the Maritime Museum of the Atlantic, detailed in this chapter. The Table, depicted in Figure 3 and Figure 5, consists of a relief wood model of the city of Halifax augmented with top-down projection mapping and augmented reality information visualized using the Microsoft HoloLens. The table's design was inspired in large part by the MIMIR prototype described in Chapter 2 and depicted in Figure 2. Due to the difficulty in acquiring accurate naval datasets and in recruiting expert participants in that domain, I used the Psychogeographer's Table as a vehicle to explore the same kinds of role-based collaboration using symmetrical and asymmetrical data representations.

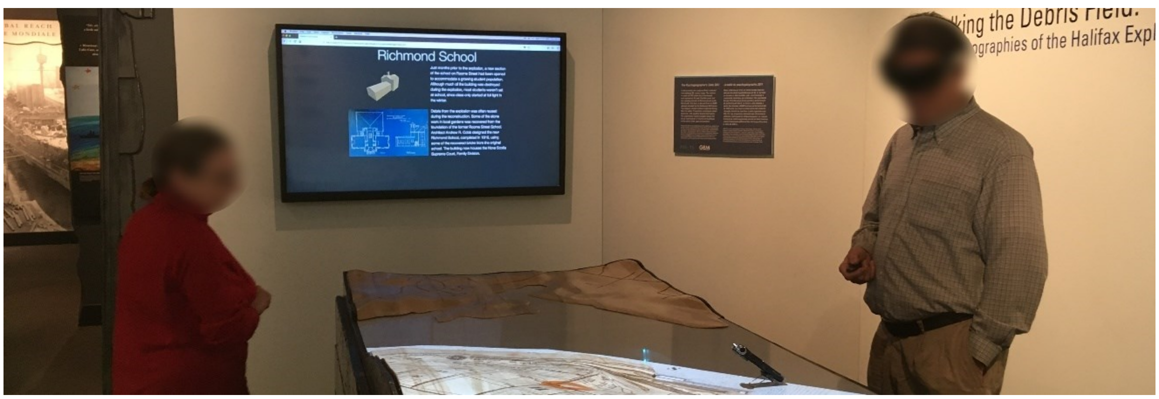

Figure 3 – Two participants interacting with the exhibit around the Psychogeographer's Table. One participant is using an AR headset rendering information about buildings affected by the Halifax Explosion, while the other participant sees a representation of the same information on the large display on the wall.

I was responsible for designing, implementing, and testing the augmented reality experience (using virtual models created by one student at the Faculty of Architecture and Urban Planning), the study design, collecting participant data, analysis, and writing. During the participant trials, I had support from two other researchers from our research lab. The study presented here was published as a journal paper on CSCW 2019 (J. Franz et al., 2019).

3.1 Introduction to the Psychogeographer's Table

Technology is ever more present in exhibits as museums look for ways to improve the visitor experience and attract the next generation of visitors. Research suggests that museums that do not invest in new technologies are considered less interesting and attract fewer visitors (Gerval & le Ru, 2015).

Museums are shared spaces, and it is essential that exhibit technologies do not create artificial barriers among visitors. People go to museums to be with other people and interact with them, have new challenging experiences, learn, and actively participate in activities (Hood, 1983).

The CSCW research community has recognized the importance of supporting collaboration in museum spaces. A common topic concerns alternatives to audio guides, moving from pre-baked solo experiences toward custom-tailored experiences using smartphones (Alvermann, 2016), tablets (Madsen & Madsen, 2015) and, more recently, augmented reality displays (Pedersen et al., 2017). Most of this research focuses on creating a common experience throughout the entire museum visit (Aoki et al., 2002; Fosh et al., 2016; Grinter et al., 2002; Tolmie et al., 2014); however, it is equally important to consider how people interact together around a single exhibit. Traditional pieces such as paintings and sculptures allow everyone in the group to see, and sometimes touch, the entire piece and thus form a shared context for discussion and interaction. I was interested in providing similar support for communication between peers when a museum piece uses immersive head-worn technologies such as virtual and augmented reality.

In this chapter's study, I proposed and evaluated three augmented reality sharing techniques designed to promote communication and collaboration in pairs while exploring a museum exhibit: Over-the-Shoulder AR, Semantic Linking, and Indicator Rings. Over-the-Shoulder AR (OtS) renders a real-time view of the AR content from the perspective of the head-worn display (HWD) on a large, shared display (Figure 4), similar to Madsen & Madsen (2015) or to commodity VR systems. To reduce latency, I rendered a virtual representation instead of displaying the view of the headset camera, which also allowed me to apply motion-stabilization prior to rendering and thus reduce discomfort for viewers. Semantic Linking (SL) uses the same display; however, instead of displaying a real-time view of the HWD, it displays contextual information about the current object selected by the HWD user (Figure 3, left). The last method, Indicator Rings (IR), is complementary to the others. It is used to indicate *where* the HWD user is currently looking and the AR "hotspots" location by rendering rings on the exhibit surface.

I was particularly interested in heterogeneous display situations where not everyone in a group of interactors had access to an HWD. I believe that this is an important scenario that needs to be addressed in museums and other domains; for example, in naval applications decision makers are often not the same people who perform data exploration however, both need to interact with data when discussing possible courses of actions and their implications. In the museum context, I observed that many visitors are not comfortable wearing a headset for a number of reasons, including fear of nausea due to previous negative experiences with HWDs, feeling embarrassed to wear the HWD in a public setting, potential impact on makeup, hair, or eyewear, and not wanting to feel obligated to engage fully with an exhibit. Furthermore, current-generation hardware for immersive augmented reality is not universally accessible. For example, it cannot be experienced by people with depth-perception disorders – even users of bi-focal correction lenses can experience issues with AR headsets. Finally, HWDs are a limited resource at exhibits, and access must be managed, especially when busy.

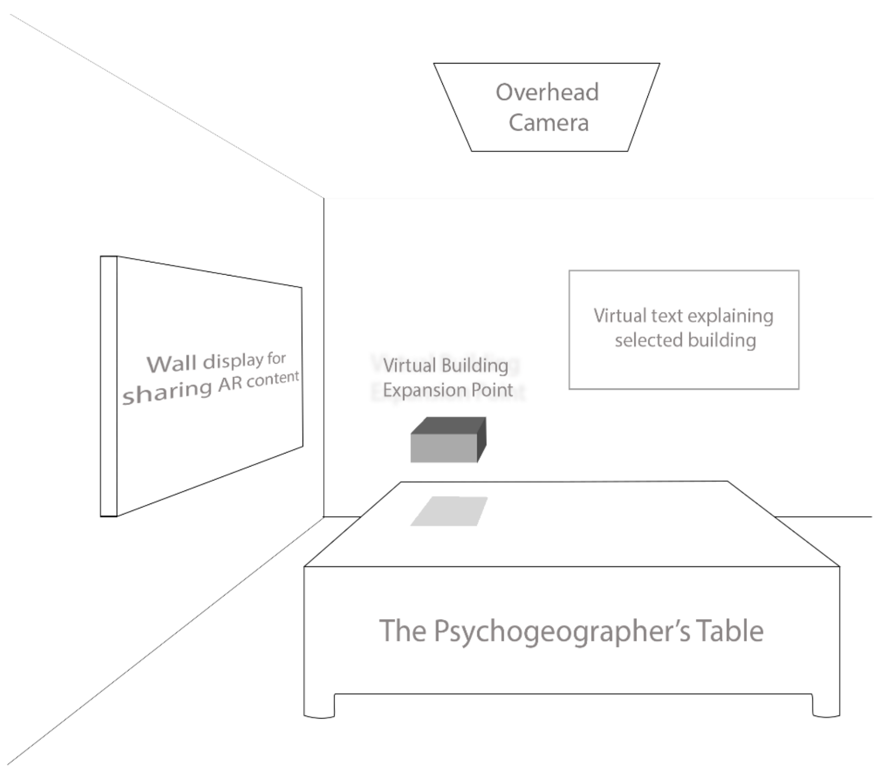

Figure 4 - Schematic of the final setup. The virtual text is rendered on the long edge wall, and a wall display is located off one short edge of the table. Virtual buildings on the table expand and

hover above the table center when selected. An overhead camera is used to record participants' movements and interactions.

With an interdisciplinary team including visual artists, an architect, and an urban anthropologist, we created an interactive exhibit that combines augmented reality elements with a large top-projected wooden relief map of the city of Halifax (Figure 3), which we refined during an in-the-wild pilot study over two months at Dalhousie's Art Gallery, and used as a platform for a four-month in-the-wild user study at the Maritime Museum of the Atlantic (MMA). In addition to collecting observations of patron activity over this period, we conducted an *in situ* between-subjects comparative evaluation of Over-the-Shoulder AR, Semantic Linking, and an HWD-only Baseline, each with and without Indicator Rings.

This work started with the technological development process needed to enable a seamless integration with the first-generation HoloLens, a projection mapping software used for the city wood model (including registration) and the integration of a large display to share real-time information. In parallel, I also worked with content development, such as image and textual information about the Halifax Explosion, and 3D modeling simplification to meet the requirements of the HoloLens mobile hardware.

3.2 Museum-Specific Related Work

3.2.1 Shared Experiences at Museums

Museology research shows that shared experiences and social interactions are essential to museum visitors (Hood, 1983). Social interaction is one of the primary determinants of satisfying museum experiences (Pekarik et al., 1999) and can promote learning (P. M. McManus, 1987) and bonding among friends, colleagues, and family (Hensel, 1987).

Importantly, visiting a museum or other cultural heritage site is not solely a social activity; individual experiences are also crucial and valuable to a visitor (Packer & Ballantyne, 2005). Sintas et al. (2014) show that social and individual museum experiences are often interwoven, while other research indicates that opportunities for social engagement and disengagement are impacted by spatial layout (Pelowski et al., 2014; Psarra, 2005) and technical supports (Sayre & Wetterlund, 2008).

3.2.2 Technology-mediated Sharing

A range of digital interventions have been employed to support and enrich the shared experience of museum visitors (Sayre & Wetterlund, 2008). In addition to typical uses of mobile devices as interactive audiovisual guides (Vlahakis et al., 2001), and for recording one's experience and sharing on social media (Weilenmann et al., 2013), research has explored collaborative authoring of exhibit narratives before, during, and after a visit (Twiss-Garrity Beth & Fisher, 2007), curated narratives that provide a shared context for exploration during a visit (Katifori et al., 2016), and mixed reality techniques for linking remote visitors to on-site visitors (Galani & Chalmers, 2013).

Some research has more explicitly explored the nature of social engagement and disengagement during the museum visit. In Sotto Voce Aoki et al. (2002), evaluated an electronic guidebook that enabled pairs of users to share or eavesdrop on content that their partner's guidebook is 'reading'. They found that mutual eavesdropping increased engagement and pair cohesion, as pairs exposed to their technique tended to stay together and maintain active verbal exchanges. Arroyo et al (2011) explore how interactive media spaces linking exhibits in different locations promote casual shared experiences among strangers. However, Hindmarsh et al. (2002) point out that more research focus has been placed on how digital technology can promote group coherence while navigating in the museum space (Fosh et al., 2016; Tolmie et al., 2014), than on how exhibits can themselves promote social interaction. Our research focuses on this latter consideration.

3.2.3 Augmented Reality in Cultural Heritage

The term "augmented reality" is used rather loosely by researchers across disciplines but generally consists of digital technology that overlays digital information on the real world via any human sensory input. Within Paul Milgram's (1994) Reality-Virtuality Continuum, Augmented reality is defined as the transition from a purely physical environment to a physical environment with digital elements added. Azuma (1997) suggests that to have believable augmented reality, we also have to guarantee correct registration of virtual objects (in the correct place, without spatial drift) and provide real-time interactivity.

For several reasons, spatially registered augmented reality (AR) is an attractive technology for cultural heritage. Curated virtual content can be placed alongside (or on) artifacts in the visitor's field of view (Chalmers & Galani, 2004; Vlahakis et al., 2001), and content can be personalized for the visitor's background, language, and interest (Pierdicca et al., 2018). While handheld AR can provide a familiar lens-like tool metaphor for exploring detail about exhibits and artifacts (van der Vaart & Damala, 2015), head-worn AR devices preserve opportunities for face-to-face communication while viewing digital (Billinghurst

& Kato, 1999) content, and are hands-free, an important consideration especially for interactive exhibits (Fleck et al., 2002).

Interest in HWDs in museums has rapidly increased with recent commercial advances, but the number of studies involving HWDs in museums is still relatively small, as a pair of recent survey papers shows. Koutsabasis (2017) reviewed 53 research papers on museum installations and found 14 that use immersive technologies such as VR and AR. However, only one used an HWD. Another survey by Pucihar et al. (2018) studied 87 applications (art installations and research studies) involving cultural heritage more broadly and found eight installations that made use of HWD-based AR, of which four were evaluated with participants (Benko et al., 2004; Cheok et al., n.d.; Herbst et al., 2008; Vlahakis et al., 2001). The work in this chapter contributes to this small body of HWD evaluations in cultural heritage and is the first study to consider HWDs in the context of shared museum experiences.

Many design considerations emerge from prior evaluations of HWD in cultural heritage applications. There are form factor considerations for the headset and any supporting devices: early work provided a compelling experience but required bulky hardware (Benko et al., 2004; Vlahakis et al., 2001). In my work, I use the relatively lightweight 1st edition Microsoft HoloLens in our research and consider short-term, exhibit-specific usage. In line with Chverst research's (2000) on locative guides, Vainstein et al. (2016) find that when navigating a site or exhibit, agency should remain with the user when selecting content to explore. The same authors recommend that physical interaction devices be preferred over in-air gestural or voice control (Vainstein et al., 2016). In their evaluation of a city-scale deployment, Herbst et al (2008) recommend that interaction mechanisms be kept very simple and that designs should permit opportunities for interacting with others. I have incorporated these recommendations in the design of our exhibit, described in the next section.

3.3 The Psychogeographer's Table

As Hornecker (2012) notes, exhibits in a museum compete for visitors' attention. Geller's work (2006) suggests that the main driver for visitors to approach an exhibit is the quality of the content. Conducting Human-Computer Interaction research in the museum context, therefore, requires compelling, polished interactive content. To that end, this work was preceded by exhibit development over a two-year period in close collaboration with artists and subject matter experts.

The Psychogeographer's Table (Figure 5) is a mixed-reality exhibit designed and built collaboratively by the Graphics and Experiential Medial Lab and the Narratives in Space and Time. The table juxtaposes a series of maps, buildings, imagery, and artifacts that evoke the complex ways in which the Halifax Explosion shaped the city of Halifax. First developed by the Situationiste Internationale (SI) in the mid-twentieth century, Psychogeography is defined by Guy Debord (1955) as "the study of the specific effects of the geographical environment, consciously organized or not, on the emotions and behaviour of individuals." The Psychogeographer's Table serves in part as a record of experiences exploring the "debris field" and its impact on urban planning and development through a series of *Derives* (reflective walks) (Debord, 1955) and Detournements (curated experiences situated in the city) (Debord, 1955) managed by Narratives in Space and Time.

A detailed account of the design and motivation of the work from the Narratives in Space and Time perspective is available in Luka et al. (2018).

The physical installation consists of a 1.2m x 2.4m computer numerical control (CNC) carved wooden tabletop representing the city's landscape and streets. The installation had a dedicated space in the museum as part of a pathway depicting the Halifax Explosion; visitors had free access to the installation space. In the center, the harbour is represented by a translucent glass surface. Under the table are two drawers containing artifacts relevant to the physical and emotional impact of the Halifax Explosion on the city and its residents that were collected on public city walks curated by Narratives in Space and Time. An overhead projector displays projection-mapped content: historical maps and current satellite images transition every 40 seconds. A large (180cm diagonal) display is affixed to the wall one meter from one of the short edges of the table and is used for the Over-the-Shoulder AR and Semantic Linking sharing techniques. Figure 4 shows a schematic of the installation.

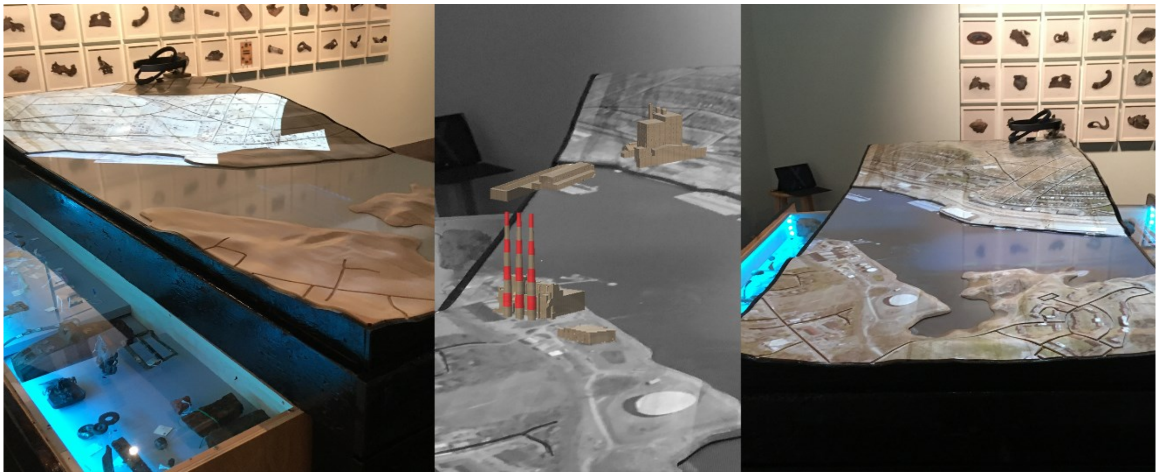

Figure 5 - The Psychogeographers' Table, a mixed reality exhibit depicting the events related to the Halifax Explosion. The centre image (black and white to enhance the contrast between the table and 3D models) shows the location of some buildings on top of the table.

In the Psychogeographer's Table, AR is used to display a set of 12 buildings from different eras of the city's history, each with unique relevance to the Halifax Explosion. Each augmented reality building (Figure 5, middle) is placed on the table surface exactly where it currently stands or stood before being destroyed.

The virtual building models are labelled with gaze-activated animated name tags (Figure 6). On the wall adjacent to the long edge of the table, I added a virtual text board that shows more information about the currently selected building. While "museums tend to draw educated visitors" (Doering, 2004), we cannot make assumptions about computer literacy (Geller, 2006) or prior experience with post-WIMP interfaces. With that in mind, I designed our interaction to be as simple as possible: the only interaction mechanisms were natural movement, gaze, and a button that participants held in one of their hands. To select a building, participants had to look, or gaze, at the building and press the button, causing the building to enlarge and hover above the middle of the table. They could then walk around the table to view the building or move their heads inside the building to see a montage of related historical and current images.

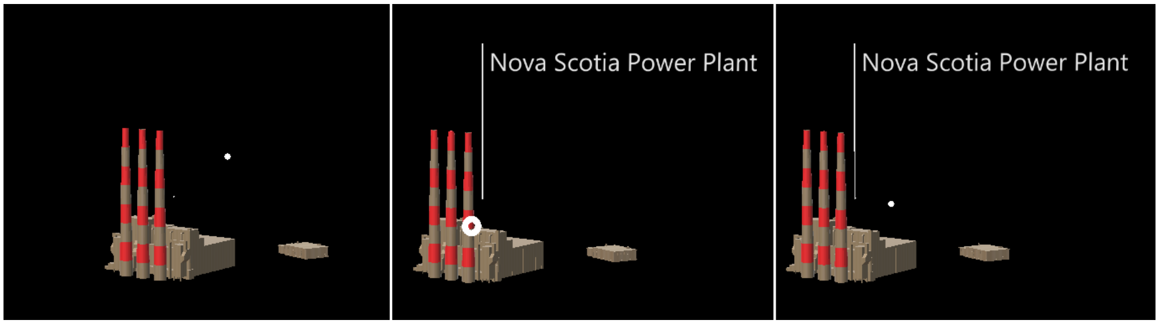

Figure 6 - Gaze-activated labels used in the final study: (left) Building models are unlabelled by default; (middle) a billboarded text label appears when the model is intersected by the HWD gaze vector; (right) Labels persist for five seconds after the gaze leaves the model. The white circle indicates the current gaze.

3.3.1 Proposed AR Sharing Techniques

I propose two sharing techniques: Semantic Linking and Over-the-Shoulder AR. I also introduce a third complementary technique called Semantic Linking, which is applied in concert with the other two.

When the HWD user selects a building, Semantic Linking (Figure 7a) displays a rearranged version of the same information visible in AR on the shared display. It shows textual information, a historical photo, and a rotating 3D view of what the model looks like in the virtual environment. Missing is the spatial context provided in AR: *where* the building is located on top of the wood model. Semantic Linking permits the no-HWD participant to have a view suitable for presentation on the wall display instead of a direct rendering of the virtual environment.

The Over-the-Shoulder AR technique (Figure 7 b) provides a direct representation of what the HWD participant is currently viewing in augmented reality, overlaid on a 3D rendering of the tabletop surface: a remote virtual scene with the same base map and buildings as those displayed in the HoloLens. The camera's location and orientation are connected to the current position and orientation of the physical headset, thus connecting both views. In this condition, the no-HWD participant can see on the wall display precisely what the other participant is observing as if looking "over their shoulder" while the HMD participant interacted with the exhibit.

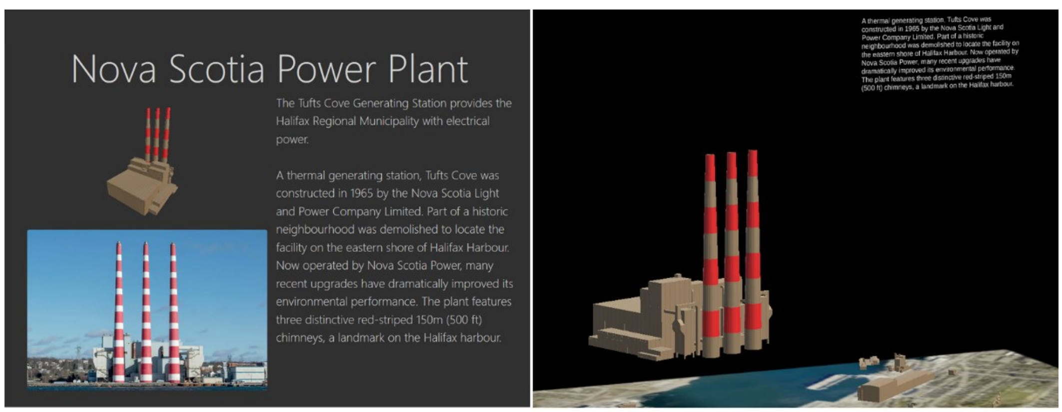

(a) (b)

Figure 7 - The two sharing methods: Semantic Linking (a), Over-the-Shoulder AR (b). Images show a screenshot from the large display used to share the AR content.

3.4 Methodology

Using The Psychogeographer's table, I explored the effectiveness of basic sharing techniques for HWD AR. I began by conducting a pilot study over two months while The Psychogeographer's Table was part of a more extensive set of exhibits commemorating the 100th anniversary of the Halifax Explosion at the University's art gallery.

In line with Geller (2006) and Hornecker (2012), we wanted our exhibit to genuinely entice visitors – who did not come to the venue to participate in a study – rather than having to persuade visitors to participate.

Other exhibit creators and I used findings from the pilot study to develop and refine the exhibit itself, our AR sharing techniques, and define the methodology for our more extensive study. The Psychogeographer's Table was installed as a feature exhibit for over four months during the tourist season at the Maritime Museum of the Atlantic. During that time, we conducted a controlled field experiment and ongoing observation of visitor interaction with the exhibit. The source code for the exhibit (HoloLens project, Top-Down Projector and Wall Display) can be accessed at: https://github.com/dal-gemlab/MaritimeMuseum-HalifaxExplosion

3.5 Museum Field Study

The field study at the museum (ethics approval on Appendix N) aimed to explore how pairs of people that know each other share and communicate when only one person has access to AR content. My goals for the exhibit were to explore shared AR experiences and investigate the following research questions:

RQ1: Does Over-the-Shoulder AR allow for no-HWD participants to share in the AR Experience? How does it affect communication and coordination?

RQ2: Does Semantic Linking allow for no-HWD participants to share in the AR Experience? How does it affect communication and coordination?

I chose pairs as a starting point as it simplified analysis and because it is common for museum visitors to arrive in pairs and small groups following the initial design of Aoki

(2002). In fact, the monthly ingress demographics for the museum while conducting the study indicated that 83% of the visitors entered the museum as individuals/small groups versus 17% as large tour groups.

I followed a between-subjects study design in the museum with a single factor with three levels: the Sharing Technique on the wall display. Participants were exposed to the following conditions: Baseline (no sharing), Over-the-Shoulder AR, and Semantic Linking (SL). SL is a middle-ground between the absence of sharing and the real-time rendering of OtS. It aligns with more traditional formats of multimedia digital content present in museums; however, its content is still directly controlled by the HWD user. I collected 30 pairs of participants over the course of three months in the museum. Conditions were alternated daily to avoid bias from visitors that might have seen the exhibit – in a different condition – before participating.

After three months, I took a brief hiatus from data collection as the exhibit was temporarily closed. I used this opportunity to reflect on my initial observations and decided to add the Indicator Rings (IR) technique for the following data collection period. Indicator Rings complement rather than replace the two sharing techniques. IR displays the *location* of the HWD participant's gazed object on the table. White indicator rings appear on the map projection when the HoloLens gaze vector intersects an AR building model on the table, gradually fading if the gaze moves elsewhere or changing to an orange ring if the model is selected. This technique can be used with or without the other sharing techniques and creates the opportunity for no-HWD participants to understand the location of virtual content without relying on the wall display.

Indicator Rings was a feature I had included in the Psychogeographer's Table design; however, I wanted to first concentrate solely on Over-the-Shoulder AR and Semantic Linking in our study. IR was added as a shared feature across all three sharing conditions, and conditions were rotated daily as before. I collected 17 more participant pairs over one month using the revised configurations, as approved by our research ethics board under an amended protocol. This extra session allowed me to consider interactions between the simple awareness mechanism that Semantic Linking affords and the sharing techniques under investigation. Specifically, I defined a third research question as follows:

RQ3: Does Semantic Linking provide a better understanding of the location of AR content? How does Semantic Linking interact with the other sharing techniques?

3.5.1.1 Participation Criteria

I recruited participants from visitors that demonstrated interest in our exhibit. To participate in our study, visitors had to approach the table in pairs. At that point, I (or the other two facilitators) ensured they were visiting the museum together, were fluent English speakers, and were at least 13 years old – this is the minimum age recommended for the Microsoft HoloLens – using the study's approved briefing script (Appendix A).

Visitors who did not qualify based on our selection criteria (except the minimum age) could still interact with the exhibit. In such cases, we took informal notes. However, we did not record video or audio, following the regulations of our research ethics committee for observing people in public spaces (Article 2.3 of the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans).

3.5.1.2 Materials

The demographic questionnaire consisted of four sections: age, gender, experience with VR (5-point), and experience with AR (5-point). The post-session questionnaire was tailored for eah participant's role. Questions are presented in Table 1. After administering the post-session questionnaire and SUS (Appendix B), I conducted a semi-structured interview with the following questions:

1. Do you think this technology can be used in other museum exhibits?
2. Were the two of you able to communicate while exploring this exhibit?
3. Were you [participant not using the HWD] able to understand what was presented in the virtual world?
4. Do you have any other thoughts or comments about your experience?

I recorded video and audio of the participants during the trials using an overhead wide-angle camera (GoPro Hero3+). I also logged positional and interaction data from the HoloLens during trials to support analysis of the number of objects selected, the number of failed attempts to select an object, movement, and time gazing at expanded buildings and the virtual text. Using video coding, I was interested in learning about differences in the amount of verbal exchange between participants, the presence/use of deictic gestures, and the level of engagement with the exhibit and their partner.

3.5.1.3 Participation Protocol

Bellow is the sequence of steps followed by the researchers on-site during the data collection period. We collected data during regular museum opening hours.

1. Each morning, researcher sets the condition of the day (alternating each day)

2. Visitors approaching the table in pairs are approached by the researcher
3. Researcher introduces the exhibit and the AR content by reading the text from the briefing script (Appendix A)
4. Researcher asks if the visitors also want to participate in the study while they are exploring the exhibit (Appendix A)
5. Participants (visitors who agreed to participate) explore the content while being recorded
6. Once done, researcher debriefs the participants (Appendix O)
7. Participants fill the post-session questionnaires (Appendices D and B)
8. Researcher starts the semi-structured interview (Appendix C)

3.5.1.4 Participation Protocol for Indicator Rings

The participation protocol did not change with the introduction of the Indicator Rings. Participants were still exposed to Baseline, Semantic Linking, or Over-the-Shoulder as the protocol mentioned in 3.4.1.3. The only difference was that Indicator Rings was activated for everyone on this stage.

Table 1 - List of Statements for Post-Session Interview

HWD	**No-HWD**
It was difficult for me to communicate with others while using the HoloLens.	It was difficult for me to communicate with the person using the HoloLens
It was easy for me to share and explain what I was seeing.	It was easy for me to understand what was being displayed in the virtual environment.
I feel that my friend/family were disconnected from me while I was exploring the virtual content.	I feel that I was disconnected from the person exploring the virtual content with the HoloLens
I enjoyed this virtual experience combined with the physical model.	I enjoyed this shared mixed-reality experience even without the headset

3.5.1.5 Data Analysis Tools

Video was coded using NVivo 12 software to calculate timing information such as the time participants spent looking at the shared wall display (a detailed description is listed on Table 2). To calculate total travel distances, I mapped the participants' positions while they were interacting with the exhibit by manually selecting their positions on the video every two seconds. I used parametric tests for the SUS scores and other Likert data according to Sullivan et. al. (Sullivan et al., 2013) and Joost et. al. (Winter et al., 2010). Due to unbalanced groups across factors, I used type-II ANOVA (car package, R 3.5.2 "Eggshell Igloo") with factors IR (on or off) and sharing method (OtS, SL, or BL). Statistical analysis used to answer the research questions presented on this study were planned before running the study.

3.5.1.6 Instruments

I used the SUS survey to measure the overall satisfaction of the HWD participant regarding the entire system, i.e., headset, table, and the shared display with the shared content. I opted only to apply the SUS to the HWD participant because I was interested in the perceived usability of someone using the HWD across the different sharing methods. Furthermore, the phrasing of the assertions in the survey refers to a "*system*" which no-HWD participant never uses and could cause confusion. The SUS score was computed as described in (Brooke, 1996).

Independent variables in the study were measured using internal sensors of the Microsoft HoloLens, audio-recorder and video coding. Table 2 illustrates the source of each variable.

Table 2 – Source of the data for the independent variables in the MMA study

Data Point	Source
Wall Time	HoloLens
Audio Excerpts	Audio Recorder Placed on the Table
Position in Space	HoloLens + Video Coding
Time Looking at TV or Wall Text	HoloLens (HWD) + Video Coding (TV, Np HWD)
Time No-HWD Looking at the Table	Video Coding
Number of Dietic Gestures	Video Coding

3.6 Results

A total of 94 visitors (47 pairs) participated in our study during both our runs at the museum; all participants completed the study. Table 3 shows the pair count per condition. Participants' age demographics are shown in Table 4. 47 participants were female, and 45 were male. I observed a curious distribution of the ages of our participants, where there seems to be a gap (bi-modal) between 25 and 40 years old. Most participants belong to the 45 to 65 years old group. Participants were generally composed of pairs of similar ages. Participants with and without the indicator rings were similar; they were tourists in the city of Halifax with a mean age of 45.45 years during the "pre" peak phase and 45.23 years in the "post" peak phase. The number of cruise ships in town was also similar: 31 vs. 27.

Experience with VR and AR trended towards the novice range (Figure 8). I also allowed visitors to experience the exhibit without participating in the study. Using our field notes, we estimated that approximately 400 museum visitors explored our exhibit with the HWD during our stay at the museum. I did not collect audio and video from non-participant visitors; however, I will share informal researcher observations regarding those users in this section.

Table 3 - Number of pairs of participants for each condition.

	w/o IR	w. IR
BL	10	6
OtS	10	5
SL	10	6

Table 4 - Age demographics of the 47 pairs of participants.

	Mean Age	SD	Max
HWD	45.38	18.65	79
No HWD	42.27	15.94	71

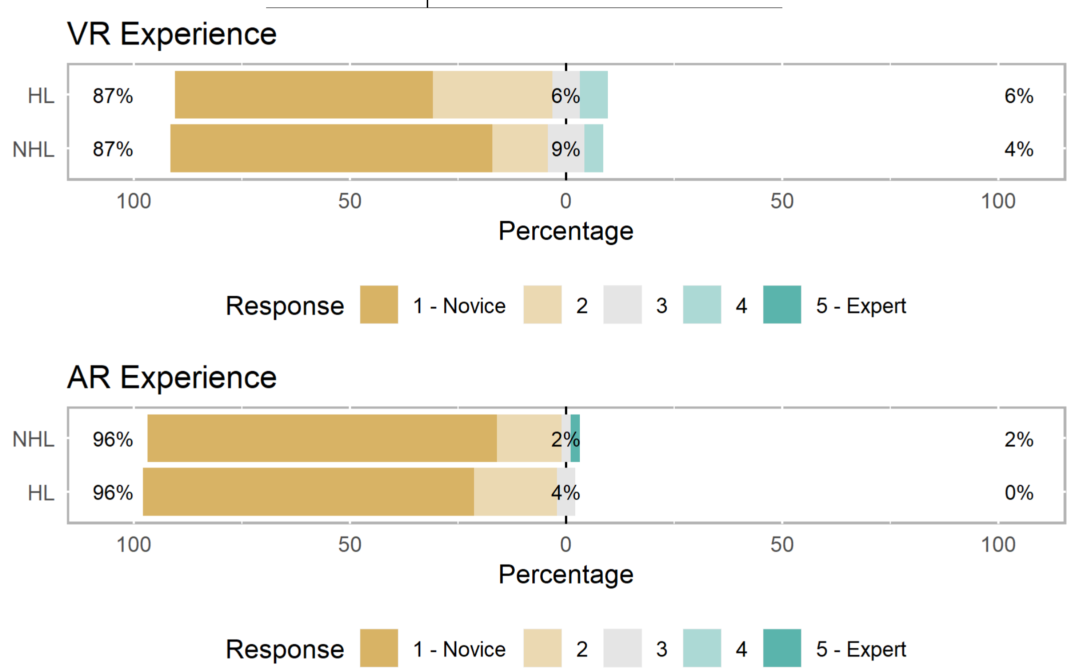

Figure 8 - Results for the self-reported identification with experience with VR and AR. Participants used a 5-point Likert scale to rank themselves from Novice to Expert. HL stands

participants with access to the HoloLens while NHL stands for participants without the HoloLens.

I informally observed in the post-session interview and after the trials (when the no-HWD display was using the HWD) that participants seemed to be rather excited and enjoyed their experience while exploring our installation with the augmented reality headset. All participants agreed with the question "*Do you think this technology can be used in other museum exhibits?*" in line with the findings of Vainstein (2016) and some even suggested different applications in other museums.

SUS scores for Baseline, Over-the-Shoulder and Semantic Linking averaged 80.78 ± 13.21, 68.83 ± 18.14, and 81.09 ± 10.16 points respectively. Data failed to reject the null hypothesis using Shapiro-wilk test $(p < 0.05)$ thus indicating a non-normal distribution. According to Norman, G (2010) ANOVA is valid even when data is not normally distributed. Two-Way ANOVA showed no significant interaction effects between the sharing methods ($p = 0.4468$) and the presence of the Indicator Rings. I found a statistical difference on the SUS scores between the sharing methods ($F(2) = 3.598$, $p = 0.0364$). Tukey post-hoc test revealed a statistical difference between the Baseline and Over-the-Shoulder conditions and between the Over-the-Shoulder and Semantic Linking conditions. There was no difference between the Baseline and Semantic Linking conditions. Figure 9 depicts the distribution of the answers.

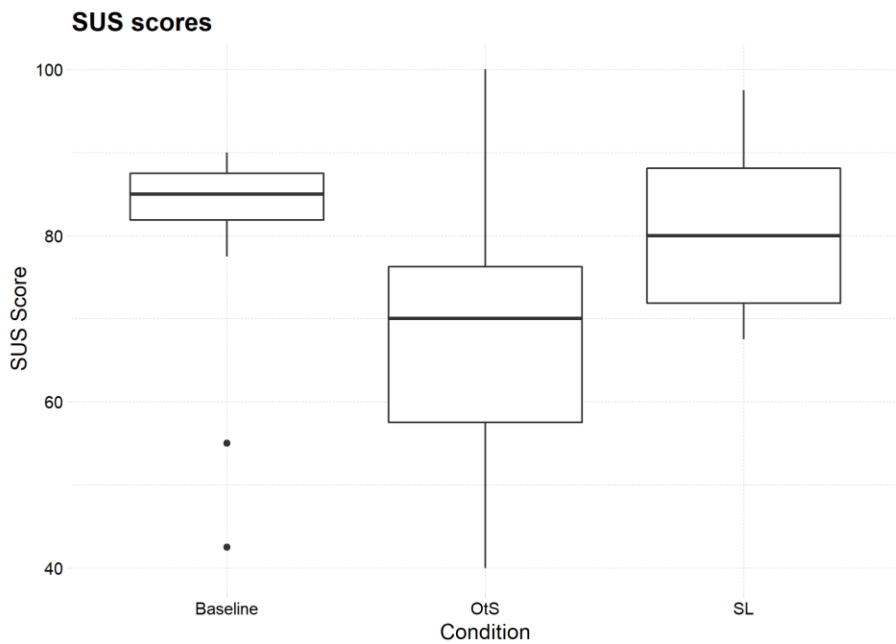

Figure 9 - HWD participants SUS scores between conditions

3.6.1 Extending AR experiences to non-AR users

Results indicated that both Over-the-Shoulder AR and Semantic Linking conditions successfully supported sharing AR content with the no-HWD participants. However, there are noticeable differences between the two sharing models.

Participants agreed that they could understand the virtual content in both sharing conditions, which was statistically different from the Baseline. However, their perception of *togetherness* or being disconnected from their partner was not equal across both sharing models. The questionnaire data shows that no-HWD participants felt statistically less disconnected from their partners when exposed to Over-the-Shoulder AR when compared to both other conditions.

Audio from the trials further enforces this difference. Participants in Over-the-Shoulder AR talked about what they saw in common:

- "What you are seeing right there [TV], is what I am seeing right here [table]."
- "Can you see what I see?" (HWD) "Yes, I can see an approximate" (no-HWD), "Look, the sugar refinery, that is something you read about it before." (no-HWD)
- "Can you see what I see? Can you read the text?" (HWD) "Yes, right here [TV]." (no-HWD)

In contrast, most of the pairs in the Baseline did not talk. The ones who did had a one-sided conversation where the HWD participant was describing or narrating what they were observing. Participants in Semantic Linking mainly discussed the textual information provided about the buildings: *"Can you tell where they are?" (HWD); "No, but I am reading here" (No-HWD)*.

It appears, however, that the sense of being disconnected from their partners was not directly related to the participants' ability to communicate with each other. Data suggests that Semantic Linking promoted a statistically higher perceived ability for the no-HWD participant to communicate with their partner. Qualitative data from the video suggests that no-HWD participants exposed to Over-the-Shoulder AR were fixated on the secondary display, probably because they wanted to gather all the possible details from the *live-streaming*, thus speaking less to their partners.

3.6.2 Bringing pairs together via shared experiences

From observer notes and analysis of the video, using a combination of the presence of deictic gestures, verbal exchanges, and dispersion of no-HWD participants away from the table, I identified three broad collaboration patterns among our participants: "independent use", "tour-guide--tourist", and "shared exploration".

Pairs in the Baseline condition explored the AR content independently from their partners. Video data shows that the small amount of conversation in this condition was mainly one-sided, where the HWD participants described what they saw.

I observed that HWD participants not exposed to Indicator Rings but exposed to a sharing method acted very similarly to a museum docent or tour guide: they controlled the exploration of the exhibit, deciding where to go and which virtual building to choose next. Without Indicator Rings, very few no-HWD participants, and then only in Over-the-Shoulder AR, asked their partner to select specific models.

Participants in Indicator Rings conditions demonstrated shared exploration. I observed from the video logs that, without the indicator rings, the no-HWD participants were anchored or fixated on the content in the shared display, especially in the Over-the-Shoulder AR condition. Observational and video data show that no-HWD participants interacted more with the table and their partners when exposed to IR. There is evidence of this behaviour from the time no-HWD spent looking at the table. Two-way ANOVA showed a significant difference ($F(1) = 21.98, p < 0.001$) with a large effect size ($\eta^2 = 0.394$) in the no-HWD participant time looking at the table variable (Figure 10), with participants exposed to the indicator rings spending more time observing the table.

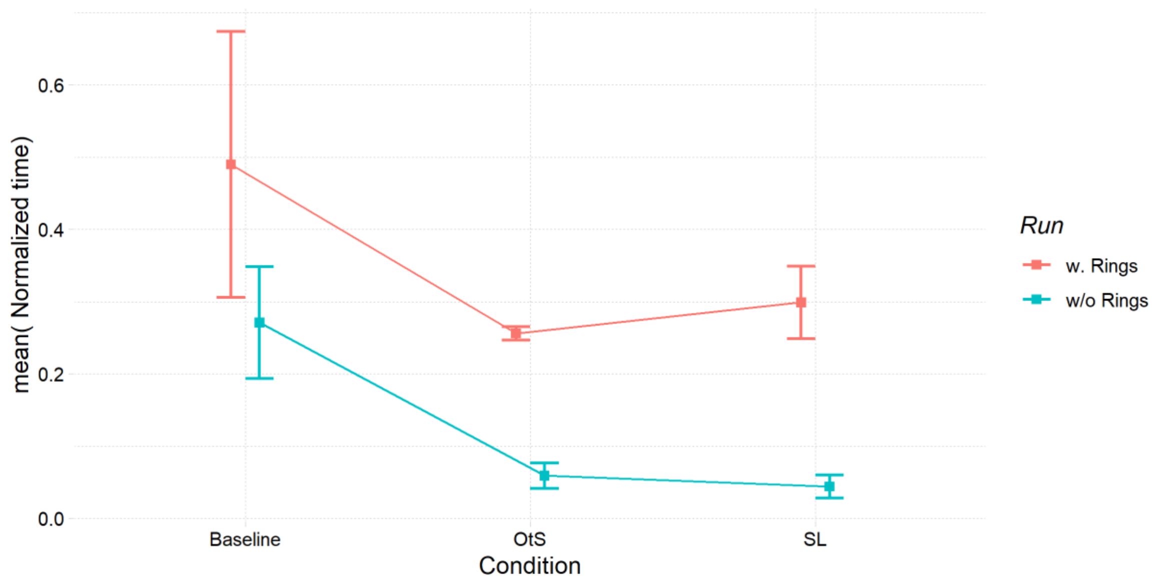

Figure 10 - Mean normalized times for time spent looking at the table by the no-HWD participant

Indicator rings allowed no-HWD participants to relate the content on the wall display to locations on the table, thus allowing them to guide the HWD participants towards specific locations/buildings without seeing the AR content. No-HWD participants frequently swapped the tour guide role with their partner using the indicator rings (Figure 11). This emergent behaviour was more prominent in the Semantic Linking condition, most likely because the wall display content contained no geospatial references, as opposed to the Over-the-Shoulder AR condition. I observed in the video analysis that HWD participants also used the indicator rings to show where buildings were on the table: "*This is where the shipyard is"* [HWD participant looks at the shipyard and uses the ring to show building location].

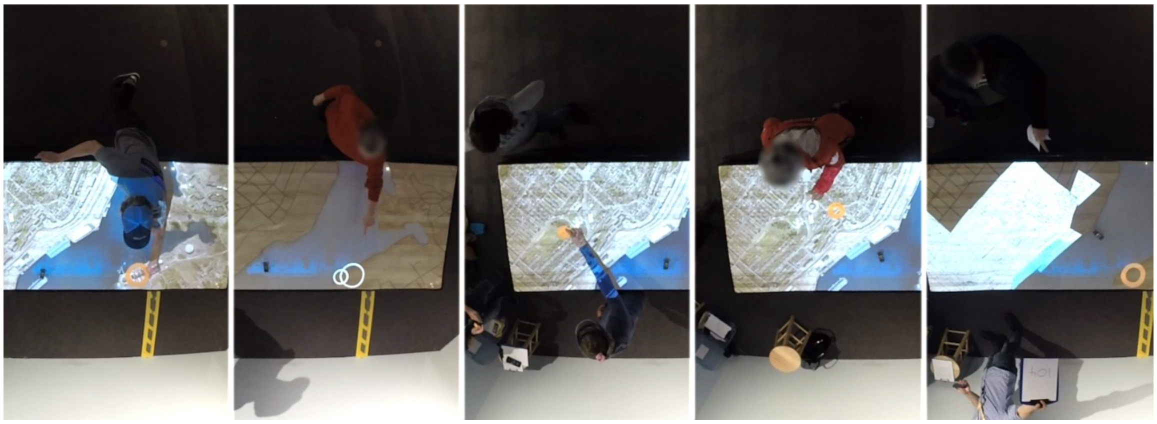

Figure 11 - No-HWD participants pointing directly at the Indicator Rings

3.6.3 Bystander Interaction

During the video analysis, I also observed the behaviour of other museum visitors, who were not participating in the study, i.e., bystanders who approached the table while participants interacted with the exhibit. The great majority of visitors approached the table, looked at the projected maps and walked away. Some visitors quickly looked toward the HWD participants but left without asking questions.

I observed, however, that the visitors who stayed more than a couple of seconds behaved differently depending on the current study condition of the day. Bystanders in the Baseline condition were attracted to a slideshow of historical photos displayed in the sharing display (the museum did not agree to leave the display off when there was no sharing). Over-the-Shoulder AR seemed to attract bystanders' curiosity towards the sharing display, especially in younger visitors; however, it appears that it is difficult for visitors to create a connection between the display and the table. On the other hand, Semantic Linking impacted bystanders in a similar way as other museum signage: visitors would stop by the display and read the information, sometimes more than one building. However, I believe that visitors treated the information on the TV as a slideshow and not something that another visitor was controlling.

Although rare, there were instances of direct interaction between the HWD participants and bystanders. I observed two instances of bystanders who made the connection between that headset and the Over-the-Shoulder AR sharing and then asked questions to the participant. There was also one occurrence in which the HWD participant started explaining to a visitor that the content she was seeing on the TV (Semantic Linking sharing) was being controlled by the participant's actions.

3.6.4 General Observations

In this subsection I introduce the aggregate of the field notes from the three researchers conducting the experiment in the Maritime Museum of the Atlantic. These notes include data from the participants during the field study and visitors (non-participants) that interacted with the exhibit but opted not to participate in the study.

Most of our participants quickly understood how to interact and select buildings using our system. The selection methodology was point-and-click, but participants had to point with their heads towards the objects rather than the selecting ray emanating from the user's hand (as in typical VR applications). "*Nose-pointing*" is a designed feature of the HoloLens 1 headset, and there is no embedded positional hand-tracking capability. Nevertheless, some participants and museum visitors seemed to want to use their hands to point at the 3D objects. After being instructed again to use their head/nose to point, some users would start pointing with their nose as instructed but keep pointing with their hands at the same time.

Because I did not block off the table area in the museum, other museum visitors would approach from time to time during study sessions. During the busy hours of the museum, it was not uncommon to have the table wholly surrounded by museum visitors. Although I opted not to run participants when the museum was extremely busy, sometimes, a wave of visitors would gather close to the table during our trials. The participants proved remarkably resilient to external interference, probably because they were so focused or immersed in the virtual content. Participants kept interacting with our exhibit even when being stared at by other visitors.

During busy times, I noticed that although many visitors were curious when someone was using the headset, very few interrupted the HWD user but instead observed the HWD user, the table, and the large display from a distance. Visitors who engaged with an HWD wearer generally had first read the exhibit description, which included a discussion of the exhibit's AR component. The exception to this rule was children, who were generally eager to ask the HWD wearer or the researchers questions about the technology.

In line with observations in prior work by Vlahakis (2001), younger visitors required almost no explanation about how to interact with the exhibit: they started exploring the AR space as soon as they put on the HWD and were given the "clicker".

3.7 Discussion

With the exhibit, I chose to explore the proposed sharing techniques in a long-term in-the-wild setting since participants behave differently when confined to trials in lab settings, especially regarding social and behavioural patterns (Hornecker & Nicol, 2012). Rather than offer extensive implications for design (Dourish, 2006), I provided a couple of observations that I believe are most pertinent. They provide a starting point for other researchers working on museum spaces to contextualize our findings to their unique circumstances and make better inferences tailored to their needs.

This in-the-wild experimentation approach is constant throughout all the studies on presented in my thesis (albeit some were adapted to follow COVID-19 protocols).

To run a in-the-wild study using a museum piece effectively, I developed a highly polished exhibit in collaboration with artists and subject matter experts over two years. One of the design goals of the exhibit was to encourage movement around the table, in part so that museum visitors would appreciate the spatially-embedded, 3D nature of the AR elements. While tabletop displays have long supported movement during collaboration, immersive displays, including most current VR headsets, tend to minimize movement to avoid people getting tangled in cables. Surprisingly participants navigated around the table with ease and were able to maneuver around other visitors without any noticeable problems, even when the space was busy. The use of HWD AR generated a novelty effect. The participants and other visitors were often amazed by the fact that they could move while the augmented buildings would stay in place (which is one of AR's constraints defined by (Azuma, 1997)):

> *"[...]But be able to walk around the cotton mill. You get more of a sense of the scale, which I like because that's always hard to get when you look at just a flat map with building locations on it".*

There was a perceived improvement in the general user experience of our exhibit from the pilot in the gallery to the final study in the museum. I posit that the increase is due to information added after the pilot study: labels and textual information about the content, in line with McManus (1989). I observed museum visitors talk about the textual information while exploring the exhibit. Adding building labels and extra information on the virtual text board increased visitor interest and autonomy since they asked fewer content-related questions of the researchers and could better explain what they were seeing to people around them.

Results from the SUS also have an impactful meaning for museum environments as it appears to be a trade-off between enabling a shared experience and users' desire for support to use the AR exhibit. It is crucial for developers and designers trying to increase engagement in pairs or groups of visitors to understand that sometimes the most direct path to sharing views might reduce engagement – not only for the current exhibit but also for future exhibits that might use the same type of technology. Moreover, museum technology should be *invisible* to visitors (Geller, 2006), or if visible as in AR experiences, it should not be intimidating as ideally, visitors should be able to interact without assistance from staff.

3.7.1 Sharing Knowledge and Avoiding Lockout

A significant concern when two or more people interact with an immersive system is that the person wearing the headset may *shut themselves off* or be *locked away* from the others because they are immersed in the virtual content. A similar process can occur for those not using the headset: they may slowly stop interacting with the HWD user as they cannot see or interact with the virtual content.

Lockout is a crucial consideration for exhibit designers and museum curators considering using immersive technologies. Traditionally visitors go to museums in groups with the expectation of fluidly engaging and disengaging with others (Aoki et al., 2002) as they flow through the galleries. I expect that given the current cost of HWD AR devices, many museums will not be able to afford many headsets, and not everyone might feel comfortable wearing one in public. Museums might also not want to have multiple headsets available to visitors at the same time in order to deter theft or damage. While HWDs in museums may be as pervasive as audio guides in the future, the sharing techniques proposed here are highly relevant since they allow people to communicate and share views of virtual content without requiring everyone to wear a headset. Moreover, I believe the techniques proposed here can be applied to other domains and visualizations that may require heterogeneous access to information. Examples include situations where decision staff might not use headsets but rely on data specialists exploring data in AR and tabletop displays with limited space around the augmented interactive surface.

I believe that two factors were in play regarding the lower dispersion – no-HWD participants moving away from our exhibit while their partner was still exploring the content – when exposed to our sharing conditions: first, with Semantic Linking and Over-the-Shoulder, no-HWD participants are given a common ground to initiate an interaction with their partner; second, no-HWD participants exposed to the Baseline condition had only one motivator to stay engaged with their partner: their curiosity to know what they saw in AR. However, I believe that HWD participants were very focused on their own exploration and required specific questions – difficult for someone who cannot see the AR content – to move them away from their own experience into a shared one. Some of the relevant comments of HWD participants in the Baseline condition were:

- "*It is hard for us to do it together when only one person is wearing the headset*" (P2, no-HWD, Baseline, no Indicator Rings)
- "*It was completely personal experience*" (P38, HWD, Baseline, no Indicator Rings)
- "*I found it because this technology is new; it felt like I had horse blinders on to some extent. I was engaged on my own, somewhat tuned out to notice who I was with.*" (P4, HWD, Baseline, no Indicator Rings)

There are distinct differences between the Over-the-Shoulder and Semantic Linking techniques that require discussion. A critical way in which Semantic Linking and Over-the-Shoulder differ concerns static vs. dynamic content since, in Semantic Linking, the content is static on the wall display as opposed to the real-time rendering in Over-the-Shoulder. This change means that when exposed to Semantic Linking, the HWD participant can freely explore the world even if the no-HWD participant is still reading the information about the last selected building. With Over-the-Shoulder, on the other hand, the HWD view is directly connected to the wall display view, meaning that the HWD participant feels obliged to wait for their partner to finish interpreting the current view before moving on. I believe that Semantic Linking can provide an advantage when people are not reading at the same pace, more likely as the number of visitors at the exhibit grows.

In that scenario, Semantic Linking allows the HWD user to "scout ahead" for the next exciting target that they want to show to the others. This static vs. dynamic difference also impacts how other visitors, or bystanders, interact with the sharing display. Semantic Linking appears to be more appropriate for busy museums because it allows bystanders to interact with the display in a format that they are more accustomed to since the display mimics traditional museum signage.

In sum, the proposed techniques effectively turned a solitary AR experience into a shared one by enabling individuals not wearing an HWD to engage with content that was once exclusive to the HWD wearer. Given the informal observations of gallery and museum visitors, I believe that both OTS and SL can be used effectively without adaptation in small groups such as families.

3.7.2 Enabling External Guidance

Over the first three months of the museum study, no-HWD participants were mainly receivers of content, that is, observers of the HWD participant's journey while they explored the exhibit. Although some no-HWD participants specifically asked for their partners to explore specific regions, these interactions were generally brief. The HWD participant describing the virtual content accounted for the small amount of conversation in the Baseline condition. Transcripts from Over-the-Shoulder showed that conversation was mostly biased towards participants agreeing that they could see similar things in the display and with the headset, followed by silence. Participants in Semantic Linking conversed about the historical details of the buildings selected by the HWD participant.

No-HWD participants in the Semantic Linking condition had no understanding of the location of virtual content on the table unless their partner explicitly showed them. Therefore, while participants enjoyed their experience, better support was needed for cooperation and coordination between our pairs. Indicator Rings supported better situational awareness for no-HWD participants, but the HWD participants also used it as a mechanism to reference regions.

Indicator rings enabled no-HWD participants to follow their partner's gaze on the table and allowed them to give directions. This new interaction is essential as secondary displays may not be present in other scenarios and, as observed from SUS scores, there may be an extra cognitive impact on the HWD user when using Over-the-Shoulder to ensure their partner "sees what they see." The presence of the indicator rings alone might not provide the same level of shared experience as when used in combination with the other sharing techniques. However, they can be enough to spark verbal exchange and casual guidance. Some HWD participants displayed buildings' locations by glancing over the buildings to make the indicator rings appear while they were saying the building names out loud. Two no-HWD participants in the Baseline condition asked: "*What's that you just selected*" when the indicator rings changed from while (gaze) to orange (selection).

As observed from the quantitative results, the presence of the Indicator Rings significantly increased the time participants spent looking at the table instead of elsewhere. Their presence significantly decreased the tendency for some participants to focus solely on the wall display, often facing away from their partner. Additionally, I observed through video analysis that the indicator rings allowed no-HWD participants to ask more questions about the content and actively take control of the exploration of the exhibit by giving instructions and directions to the participant wearing the headset.

3.7.3 Impact of Sharing on Usability

There was an increase in the SUS scores (higher is better) compared to the pilot study: from 66.56 to 80.78. Using the SUS score as an indicator it appears that the design changes we made in response to the pilot improved the usability

I posit, however, that the presence of the real-time rendering negatively impacted the system's usability as perceived by the HWD participants. Although HWD participants felt the least disconnected from their partners (post-session questionnaire assertion 2) in the Over-the-Shoulder AR condition, I posit that constantly controlling what is visible on the shared display can constrain coordination and can be taxing for HWD participants.

3.8 Conclusion

The end of this chapter marks the transition of my work from Mixed Reality group work into computer supported group sport work, more specifically road cycling. Enabling collaborative interaction with museum exhibits is important as museums are usually explored in groups of two or more people. Augmented reality is a great asset to museum curators since it can be used not only to create new exhibits from scratch, but it can also be used to add extra layers of content to current exhibits such as paintings and sculptures or to replace traditional audio guides with more immersive experiences. AR headsets, however, are still considered a novelty and museums might not have the budget to purchase multiple devices. Hence, we must seek ways to enable visitors to interact with engaging mixed-reality exhibits even when not all members of the party have access to augmentation devices.

I showed that both of our techniques made it possible for visitors without the headset to understand what was presented in the virtual part of the Psychogeographer's table. I also showed that our participants felt that they could communicate with each other when using our sharing methods which is most likely a consequence of their enhanced comprehension of the exhibit content. This is an important finding for future group work, including the rest of my thesis with cyclists, as it demonstrates that sharing a full representation of the information with others is not required to promote grounding for conversation and engagement.

Furthermore, I showed that one could add an extra augmentation layer to the shared projected content on the exhibit in the form of indicator rings revealing the gaze and selection activity of the headset user. By doing so, I observed a significant increase in the interaction between participants and noted that the presence of the indicator rings provided enough spatial information to participants without a headset to enable them to guide the headset user.

3.8.1 Implications for Cycling Group Work

Results from this work inspired the ideation process for the cycling work that follows. Riders are similar to participants in the baseline condition; although each has their individual data collection and visualization tools, their data is not shared across riders. They work together, but information that could help enhance work is kept in knowledge silos.

Approaches like Semantic Linking and Over-the-Shoulder could be adapted to support the riders' group work. Since riders have limited screen space in their devices, asymmetrical information sharing such as SL is an interesting approach, especially since one-to-one sharing was not required to promote engagement. Furthermore, supporting the role-based behaviour observed in SL is vital to the following cycling work. Using SL, participants have demonstrated the ability to work together while impersonating two, albeit informal, roles, which is expected from group cyclists as they too impersonate different roles as they ride together.

3.9 Limitations

Any instrumentation induces changes as the instrument itself modifies the subject. In the museum the other researchers and I were the instruments studying and possibly modifying the participants behaviour from their interaction with us, questions and answers, or just by us being there observing as they wandered around our exhibit. While we tried to not intervene while participants were exploring the exhibit, our presence there could have limited self-exploration with the content and the technology (e.g., trying to figure out how something works vs. asking questions to us).

While we tried to *recruit* all pairs that showed interest in our exhibit, some people might have felt uncomfortable with us asking them to participate in the study either due to perceived time requirements or due to socio-cultural reasons from their upbringing (for example we only had male researchers conducting the experiment). We were also limiting entities for participants coming from abroad (e.g., visitors coming from Europe on cruises) that did not speak English fluently as it was the language we accepted in the study because it is the only common language we understand as group. Furthermore, we might have created participant bias on what to expect from the technology, interaction techniques and content presentation as we had to introduce the study's technology using our perspective of it and domain expertise.

Analysis was done by the same researchers that conducted the experiment which can bring pre-defined ideas from the study sessions. Lastly, some behavioural patterns that were not related to collaboration and interaction between the pairs and exhibit might have been missed as the researchers conducting the analysis were focused on that topic primarily.

CHAPTER 4 CYCLING TOOLS AND NOMENCLATURE

In the following chapters, I switch gears to collaboration among road cyclists. My previous work on supporting people with different access to information or having different roles is a cornerstone in inspiring the rest of my thesis. Road cyclists are much like the museum visitors discussed before. They work together with a shared objective while performing different tasks to support the group. Although they are trying to collaborate, they do not have tools to support access to the information collected by their peers while riding, thus creating silos of knowledge hindering collaborative work.

This chapter is an opportunity for the reader to become familiar with the nomenclature and technology present in modern bicycles if they are not aware of the current state of cycling. Readers familiar with the sport are welcome to skip this chapter in its entirety.

While riding, cyclists rely on their handlebar-mounted computer, known as a cycle computer, cyclometer, or bike computer, as a mechanical device attached to the bicycle's front wheel that could inform the cyclist how far they have ridden.

Figure 12 - Veeder's Cyclometer. Source: Canadian Cycling Magazine, Dec 2018 - Jan 2019

Modern bike computers use digital displays and are separated into at least two components. A head unit displays the information collected by sensors or calculated internally, such as the average speed and a wheel magnet paired with a reed sensor attached to the bike's fork to collect speed. Basic computers, such as those depicted in Figure 13, are popular with casual riders, like commuters, due to their low cost and ease of use. Most measure speed, distance, and time, while some also measure cadence (how fast the crank arm[3] is spinning) via an extra sensor on the bicycle's crankset. In the remainder of this chapter, I will introduce more advanced cyclometers used by enthusiasts, serious athletes, and professional riders. I will also present data sensors and mobile applications that can be employed instead of the head unit or used for post-hoc ride data analysis. Appendix O contains an excerpt from "Ride on" Magazine naming most of the parts of a road bicycle in case the reader is unfamiliar with the terms.

Figure 13 - Cateye Strada Slim basic cyclometer. Source: https://www.cateye.com/

4.1 Smart(er) head units

During the mid-2000s, cyclometers started to adopt GPS tracking, mapping features and a new Adaptive Network Topology (ANT+) (*The Wireless Sensor Network Solution -*

[3] The reader can find a diagram of a bicycle's parts and names on Appendix O

THIS IS ANT, n.d.) sensor network used to connect multiple data sources. While these new devices are more capable than older generations, how they deliver information to cyclists is still reminiscent of the original digital head unit. Most of the instantaneous information is assembled in a grid-like pattern (personalized by each cyclist) which can occupy one or more "pages" (Figure 14). Spatial information such as maps, routes, and altimetry have their own pages (Figure 15).

Design and information delivery are solely focused on the rider’s sensors as there is no sharing of information between multiple head units riding in the same pack, even though both ANT+ and Bluetooth Low Energy (BLE) could support such exchange. Sensors paired in the ANT+ network can transmit to multiple units simultaneously (e.g., a head unit plus a regular computer during an indoor training session). The units can sense the environment but are not trying to infer context from it.

Figure 14 - Personalized data fields for several models of head units. Although units are from different manufacturers (Garmin, Wahoo, and Stages), their layout is similar. Source: (www.garmin.com and ca.wahoofitness.com)

Figure 15 - Map views of two different head units. Garmin Edge 530 (left) and Wahoo Roam (right) Source: www.garmin.com and ca.wahoofitness.com

Like driving a car, riders are not constantly looking at the data fields on their cyclometers. Consulting information means losing focus on the road ahead, which could cause an accident. Although most of the vital data is aggregated (by the rider's preference) in one data page, switching between pages (from data grid to map or cue sheet) takes time and requires focus.

Modern head units are connected to the internet via the rider's phone, and the computers even use some cloud services such as Strava's Live Segments (*Live Segments – Strava Support*, n.d.) or the Beacon Safety Feature (*Strava Beacons*, n.d.). However, they do not leverage the smartphone or cloud-computing platforms to provide context-based layouts, which could reduce the need for pages and assist riders to focus mainly on the road.

4.2 Cycling sensors

A modern road bicycle is much more instrumented than Veeder's original sales pitch: *"It is Nice to Know How Far You Go."* Both enthusiasts and professional riders have sensors on their bodies and on their bicycles. This section will present a list of typical sensor units used in road cycling.

4.2.1 Speed

Distance (and thus speed) is one of the first metrics implemented in cycling computers. Modern computers sense speed in two ways: GPS and dedicated Speed Sensors. The latter comprises either a sensing unit (using ANT+ or BLE) attached to the frame of the bicycle and a magnet on the wheel or a sensing unit attached to the wheel's hub. The system is configured with the wheel's circumference, and speed is inferred from the wheel's rotation per minute (RPM) and angular velocity.

4.2.2 Cadence

Cadence is defined as the rotation rate of the bicycle's crankset arms, i.e., how fast are the rider's feet spinning the pedals. It is a fundamental metric in cycling related to muscle activation and active resting (Macintosh et al., 2000). A well-trained rider will control their current cadence (by switching gears or stance) depending on what they are currently doing on the road and how tired or rested they are.

Cadence sensors are very similar to speed sensors. They are composed of a magnet attached to one of the crankset arms and a magnetic sensor attached to one of the chainstays.

4.2.3 Elevation (and Grade)

Elevation (vertical distance from the ocean level) is primarily measured as a cumulative delta of the starting elevation of any given ride. Cyclists use the accumulated elevation from the start of the ride to gauge their cumulative effort. Accumulated elevation can also be used before a ride (from GPS elevation of a planned route) to have an idea of the difficulty of the ride.

Grade, on the other hand, indicates how steep any climb is during the rides. This information is an important metric, as it can be used to determine how challenging a climb will be. The grade is also calculated on the go by cycle computers, and riders can adjust their pace and cycling stance on any given climb by consulting the real-time grade.

4.2.4 Power (Watts)

Power is another important metric in cycling, albeit novices do not widely use it as the sensors (power meters) cost is high. Power meters can measure the power output of a cyclist at five different regions of the bicycle: Wheel Hub, Pedals, Chainrings, Crank Arm, and Bottom Bracket. Most meters use strain gauges to measure the deformation of the material when exposed to the cyclist's leg power.

Because power is a direct measure of effort and energy expenditure, cyclists can match their current power output to their Functional Threshold Power (FTP) while training. The ratio between current power and FTP defines seven power-based training zones (Table 4) with different physiological responses.

Table 5 - Power-based training zones as a function of FTP

Zone	%FTP	Name
1	<55%	Active Recovery
2	56-75%	Endurance
3	76-90%	Tempo
4	91-105%	Threshold
5	106-120%	Anaerobic
6	>121%	VO_2 Max
7	N/A (Short Burst Max Effort)	Neuromuscular

4.2.5 Heart Rate

Heart rate sensors are commonly used across multiple sports. Typically, the sensors are sold in three variations: chest band, wrist band, and arm band. These sensors are cheaper than power meters and, although not as precise, can be used to estimate effort. Heart Rate-based training zones can be either estimated via a formula based on the maximum heart rate (which needs to be tested or estimated) or created from the athletes' ventilatory thresholds via a VO_2 max test. The use of heart rate sensors is mostly uniform among riders, differently from power meters.

4.2.6 Rearview Radar

Relatively new to the cycling athletes, rearview radars such as the Varia RTL515[4] play a different role than previously listed sensors. According to Statistics Canada (*Circumstances Surrounding Cycling Fatalities in Canada, 2006 to 2017*, n.d.), 73% of total fatalities involving cyclists between 2006 and 2017 were caused by motorized vehicle collisions. Radar sensors give cyclists an early warning of incoming vehicles via visual and audible alerts from the head units with which they are paired. Information includes the arrival time and relative speed of the incoming object (car, motorcycle, train, elephant), and it is encoded in two levels: normal overtake or high-speed overtake.

These sensors are the first ones to try to address issues regarding the safety of the athlete instead of their performance. Radars report on an individual setting, that is, only to the head unit that it is currently connected. Grouping multiple radar units to create a larger sensing area is currently not done in cycling.

4.3 Data visualization tools

While riding, cyclists focus on instantaneous data like speed, cadence, power, and grade, or near-instantaneous data such as 3-sec/10-sec average power or average speed. Physical screen real estate and safety concerns are two limiting factors for temporal data visualization during the exercise. Temporal information is visualized before or after the exercise using tools such as Strava[5], Ride with GPS[6], and Garmin Connect[7]. In this section, I will present some of the tools provided by Strava as it is, to date, the most popular

[4] https://buy.garmin.com/en-CA/CA/p/698001/pn/010-02376-00
[5] https://www.strava.com
[6] https://ridewithgps.com/
[7] https://connect.garmin.com/

application[8] (the reader can think of as both an analytics tool and specialized social media) among enthusiasts and professional cyclists. Furthermore, Strava was the tool used to promote recall during the Remote Inquiry interviews presented in the next chapter.

4.3.1 Individual Statistics

Like on-the-ride data visualization, most post-hoc data analysis tools focus on the individual. In this subsection, I will use a ride[9] from the professional cyclist Laurens ten Dam (aka LTD) to present some visualization features available on Strava's website. The *Overview page* (Figure 16) presents general information, map, altimetry, and social features of the platform, such as who participated in this ride and the Strava segments. Furthermore, this screen shows weather information such as temperature, feels-like, wind direction and speed for the owner of the ride and their friends.

The *Analysis page* (Figure 17) displays three coordinated views: a map, an altimetry view, and the temporal sensor logs from the ride. This page provides the tools needed to explore information like power while relating it to the ride's map. Riders can explore the ride as a whole or select any period for a summarized view. Furthermore, this screen also displays any Strava segments in the route as horizontal bars underneath the altimetry map.

[8] https://www.pcmag.com/reviews/strava
[9] https://www.strava.com/activities/2767675842/overview

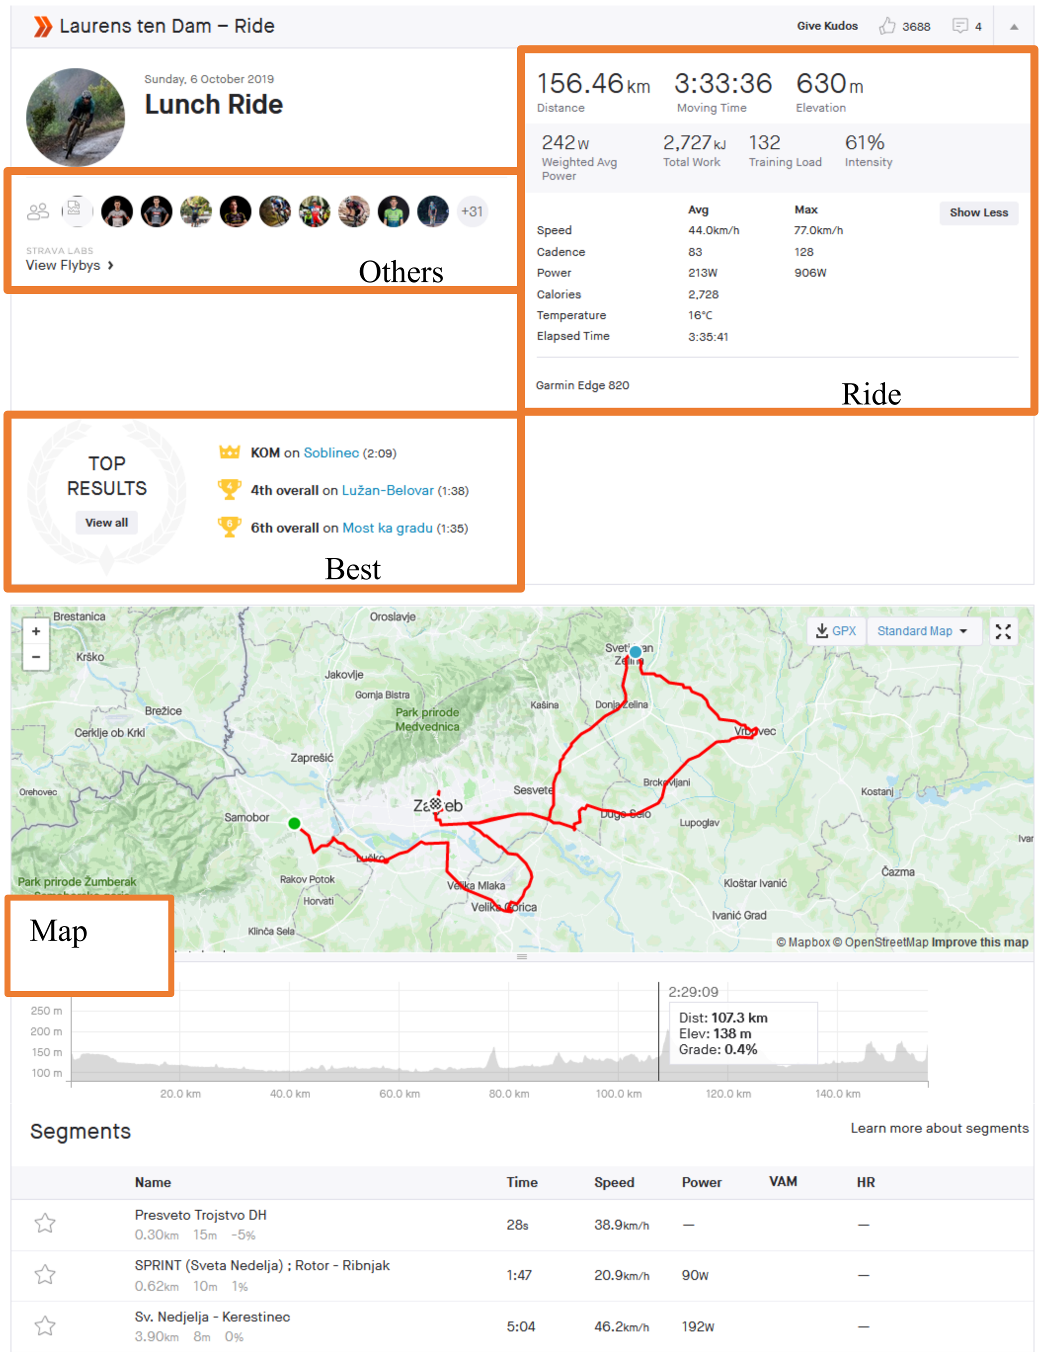

Figure 16 - Screenshot of the Summary view of LTD's 2019/10/06 ride. The screen shows a ride summary, list of other riders that were present, the ride's map and segments.

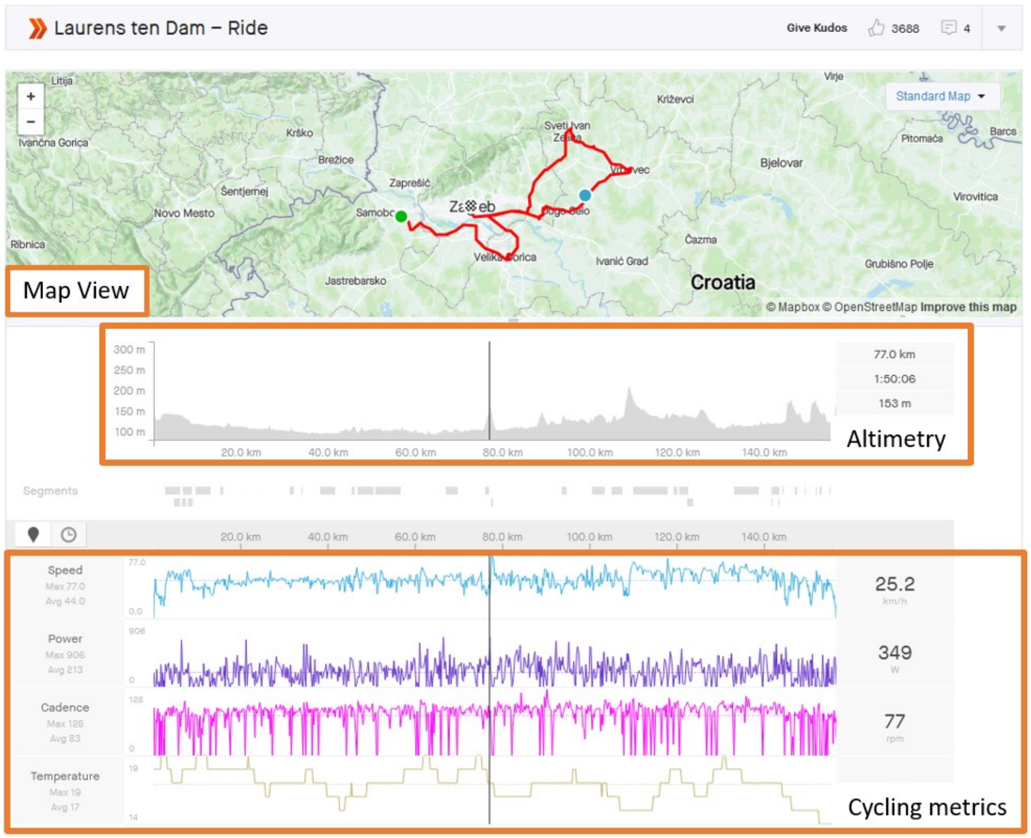

Figure 17 - Screenshot of the Analysis view of LTD's 2019/10/06 ride. It maintains the map view and presents a detailed time-series view of the ride's data. It allows for synchronized scrolling of all three regions.

CHAPTER 5 EXPLORING GROUP ROAD CYCLING THROUGH CONTEXTUAL DESIGN

The practice of sports is a crucial part of a healthy life for many. Physical activity promotes physical fitness (Davidson & McNaughton, 2000) and contributes to a healthier social life (Fisken et al., 2016). The practice of team/group sports can create a unique sense of community among athletes by creating social spaces and bringing together people who share common interest (Warner et al., 2012).

Both professional and amateur athletes track a range of data from body-worn sensors and external sources. These data can also promote socialization via sport-specific social media platforms and their community challenges (e.g., Strava's cycling and running challenges (*Strava Challenges – Strava Support*, n.d.)). Moreover, sports data can be used by professionals and enthusiasts to analyze performance and explore areas for improvement. However, data is mainly used for post-hoc analysis, for example, in tennis (Polk et al., 2014), soccer (Janetzko et al., 2015), baseball (Lage et al., 2016), hockey (Pileggi et al., 2012), and climbing (Niederer et al., 2016).

Websites such as Strava, RunKeeper, and MapMyRide, combined with low-cost sensing tools such as running watches and GPS-enabled bike computers, are contributing to the popularization of post-hoc running and cycling data visualization. However, athletes in such sports can also receive some data while engaging in physical activity via cycling computers[10] or running watches. Still, there is a lack of commercial tools for sharing data during the activity, which could provide valuable context-specific insights for the

[10] Although some cyclists, like commuters or novice riders, might use smartphone applications while they ride, they are not in scope of this research.

individual and a group of athletes exercising together. Similarly, while there has been a significant increase in publications related to sports visualization (Perin et al., 2018), most focus on post-hoc analysis (Gudmundsson & Horton, 2017) and few explore the delivery of information directly to athletes during the practice of sports (Kiss et al., 2019).

Cyclists, for example, have access to most personal data collected by body-worn or bicycle-mounted devices via small heads-down displays attached to their handlebars, known as cycling computers. Cycling in most variants, however, is a team sport where a group of riders collaborate towards a common goal by rotating the position of riders on the peloton to share the load of pushing the group. The objective varies from finishing first in a race segment (like the Giro d'Italia (*Giro d'Italia 2022 | Official Site*, n.d.) and Tour de France (*Official Website of Tour de France 2022*, n.d.), for example) to completing a 3-day long-distance challenge such as the Paris-Brest-Paris event. Despite being a group sport where riders collaborate with each other while riding, there is very little research on visualization tools that provide in-situ group information among riders in a peloton to support sharing of their data with other riders. Although cyclists need to collaborate among themselves towards an objective, their tools of the trade are individualized, thus creating small silos of knowledge instead of enhancing group-wide cognition.

Understanding the sport and the context in which athletes are immersed while exercising is essential. While post-hoc visualization tools exist, they are situated in a different contextual realm and cannot be directly used in situ. We can draw a similar conclusion regarding the use of existing in situ individual visualizations: they are not enough to promote or contribute to a distributed knowledge of the team. To develop new tools, we need first to understand the work done by athletes and the context in which this work is done. Per Perin (2018): "The process of designing visualization and interaction techniques for sports requires a deep understanding of sports."

In the study presented in this chapter, I consider cyclists engaged in the sport (i.e., not commuters) and have access to their personal data collected by body-worn and bike-mounted devices via their bike computers.

A modern, fully outfitted road bicycle can collect speed, cadence, power (the value, in Watts, of how much power the cyclist is transferring from his body to the bicycle), heart rate, road incline, pedal power phases, and even a rear radar view of incoming traffic. All this information is available to the cyclist during their ride. However, the current state-of-the-art devices rely on dashboard-like visualizations, have small displays, and inherently force the rider to look down and away from the road. Furthermore, implicit information, such as the Rate of Perceived Effort (RPE) (BORG, 1982), cannot be sensed like other ride data. Moreover, while some absolute measurements only make sense to the rider, other relative measurements like the Percentage of the Functional Threshold Power (%FTP) (Borszcz et al., 2018), current Watts per kilo[11] (W/kg) and RPE could be used by the group to make changes in intensity and leader rotation times during the ride to maximize riding efficacy. Nevertheless, no research indicates how the information could be captured and shared among riders.

[11] The current power output normalized by the rider's weight (power to weight ratio) without considering the bicycle's mass. This metric is useful to determined the rider's overall strength in climbs.

To develop new tools, we must first understand the work done by athletes and the context in which this work is done (Perin et al., 2018). In the study presented in this chapter, I explore the following research question:

RQ1. *What are the communication and collaboration challenges found by group cyclists, and where does technology fail to support them?*

I will present the results of a Contextual Inquiry (CI) study that focused on understanding the work done by groups of road cyclists. I conducted the study with three cycling groups using a modified CI protocol adapted to the COVID-19 pandemic. Results are presented using the five work models proposed by Contextual Design (Beyer & Holtzblatt, 1997):

- **Physical Model:** shows how people organize their environments to make their work easier.
- **Flow Model:** reveals the formal and informal workgroups and communication patterns critical to doing the work.
- **Cultural Model:** shows how people are constrained and how they work around those constraints to make sure the work is done.
- **Sequence Model:** shows the detailed steps performed to accomplish each task important to the work.
- **Artifact Model:** shows the artifacts that are created and used in doing the work.

The study's ethics approval can be found in Appendix J.

My work provides an in-depth understanding of the collaborative work done by groups of road cyclists, that is, riding as a group while sharing the effort towards a common goal e.g., going faster or longer. I include communication strategies, behavioural patterns, and attrition points when riding as a group and expose potential research directions for interactive systems to support group riding. Lastly, I use the data to reflect on how riders (and potentially athletes from other group sports) could benefit from tool support for real-time collaboration when they are engaged in their sport, using group pacing as a starting point. The work presented on this chapter was published on DIS 2022 (J. Franz & Reilly, 2022)

5.1 Data in Road Cycling

A modern, fully sensed road bicycle can collect speed, cadence, power (the value, in Watts, of how much power the cyclist is transferring from his body to the bicycle), heart rate, road incline, pedal power phases and even a rear radar view of incoming traffic. All this information is available to the cyclist during their ride. However, the current state-of-the-art devices rely on dashboard-like visualizations, have small displays, and inherently force the rider to look down and away from the road. Furthermore, implicit information, such as the Rate of Perceived Exertion (RPE) (BORG, 1982), cannot be saved along with the ride data. Moreover, while some absolute measurements only make sense to the rider, other relative measurements like the percentage of the Functional Threshold Power (%FTP) (Borszcz et al., 2018) and RPE could be used by group riders to make changes in intensity and leader rotation times during the ride to maximize riding efficacy. Nevertheless, no research indicates what these visualizations should look like and how the information could be shared among riders.

Communication bandwidth among riders is relatively low as current digital devices offer no communication strategies. Cyclists rely primarily on a pre-defined set of hand gestures to inform of road dangers, such as holes, rotation intervals, and other ride-related details. This communication modality suffers from high latency, requires line-of-sight, and often results in package drop in fast high-risk interactions. For example, one cyclist signals something close, but the second or third cyclist fails to transmit the information further down the line, and someone hits a pothole, tree branch or a downed cyclist.

Riders also communicate using verbal utterances. Similar to gestures, bandwidth and range are relatively low. Spoken word also requires message re-dispatch, e.g., the last rider hears a car approaching from behind; they scream "*Car back*"; all cyclists forward will re-dispatch the message (unless it is dropped): "*Car back.*"

From this perspective, we can portray road cyclists as people with situationally-induced impairments and disabilities (SIID) (Sears et al., 2003). While performing their activity, they are subject to sensory disruptions, mimicking someone with disabilities. Riders are not only subject to wind noise (mimic hearing impairment) and the fact that they are always mobile but depending on the kind of ride they are in, they might also be susceptible to long periods of exposure to high temperatures under the sun, rain, fog, extreme fatigue (some rides can last up to 3 days with short breaks), and stress due to unsafe road conditions.

Group riders are, in essence, similar to Pascoe's (2000) fieldworkers using a mobile computer system while on the go: "We identify four specific characteristics of this class of users: dynamic user configuration, limited attention capacity, high-speed interaction, and context-dependency." They have specific roles (albeit very fluid, as discussed in the following chapters) and must perform specific tasks and pay attention to a subset of information based on each role. The information must be delivered in a format that can be easily read in a short period while the cyclists (try to) maintain focus on the primary task of riding.

Another challenge in data delivery for road cyclists is the size and low resolution of bicycle computers. For example, Garmin's flagship bike computer, the Edge 830, has a resolution of only 246 x 322 pixels. Moreover, interacting with such devices can be tricky sometimes. It can create a hazard for the cyclists since they must momentarily lose focus of the road to gather information from the device and might have to remove one of their hands from the handlebar to touch the device. Future tools to support group cycling should consider these situations when defining the context in which cyclists are immersed.

5.2 Sport Specific Related Work

We are surrounded by digital data collection in our lives, from our computers' telemetry data to our purchase habits on our customer fidelity cards. It is only natural that we can also collect and track data related to our sports activities. Sports metric data collection and feedback methods have been explored in the past (Aranki et al., 2017; Baca et al., 2009; Fister & Fister Jr, 2012; Savage, 2010; Shiro, 2019); the same is valid for individuals (Khot et al., 2020; Polk et al., 2014) and team visualization tools (Perin et al., 2013; Probst et al., 2018).

According to Huang et al. (2015), our personal data "has enormous potential […] to make positive changes in our personal lives […]". This is especially true when visualization tools provide insight for users with little to no experience with visual analytics. The sports field is rather interesting in this aspect: it is composed not only of highly technical experts in elite teams, but also of sports novices. The former is collecting data from the beginning of their careers and trying to make sense of it all to promote self-improvement. While in some sports, experts and novices will not share the same space and tools (e.g., volleyball, basketball, and soccer), there are others, such as running and cycling, where professionals and novices are more likely to share the physical space (like roads), data sensing tools like running and cycling computers, and data visualization platforms.

5.2.1 Sports Visualization

In a recent state-of-the-art review on sports visualization, Perin et al. (2018) analyzed 98 peer-reviewed articles from researchers and practitioners on sports. While most research papers focused on new visualization techniques (especially for team sports), none investigated athletes' data visualization while exercising. Furthermore, the state-of-the-art review paper by Perin et al. shows that cycling visualization is under-represented with only four papers and running (cycling's closest counterpart related to data collection) also has only four papers. Most works in the sports visualization fields tend to focus on team sports where players are engaged in a limited area or region of interest, such as soccer (Perin et al., 2013), baseball (Lage et al., 2016), and basketball (Losada et al., 2016). Having all players in a field or court creates a well-defined frame of reference, facilitating data collection and analysis. Furthermore, these "court-based" sports are traditionally recorded using cameras from several angles during the matches, thus providing a rich data source for computer vision and machine learning algorithms.

As expected, most sports visualization research focuses on generating and understating game statistics to aid players and coaches in improving their performance for the next match. Current work shows a trend towards the post-hoc analysis of game data. Unfortunately, due to the particularities of each sport, there is no silver bullet solution that applies to multiple sports simultaneously, leading to fragmentation of research and findings among niche sub-fields. For example, Wu et al. (2018) proposed a framework to analyze table tennis data from multiple tournament matches, while Polk et al. (2014) worked on a similar visualization technique for tennis. Even though both sports are similar, there is not a single solution that fits them due to their unique characteristics and different visualization needs.

Not all research focuses on post-hoc analysis of gameplay data, however. Legg et al. (2012) created a glyph-based visualization system to help rugby coaches make "mission-critical decisions" while still paying attention to the current match details. Crowell et al. (2010) tested (although in a lab environment) if runners could use real-time feedback of their pace and ground reaction forces to reduce impact during the run activity to reduce the chance of impact fractures. Colley et al. (2018) proposed a simple built-in shoe display to assist runners in reaching their target pace while practicing the sport. In Clairbuoyance (Kiss et al., 2019), the authors focused on providing real-time feedback on swim direction for open water swimmers.

5.2.1.1 Cycling Specific: Pre and Post hoc

Specific to cycling, pre- or post-hoc data visualization research can be mainly divided into two areas: visual analytic tools to help understand general cycling behaviours (such as for city planning) or the visualization of one's exercise data to help learn and improve in the sport.

For example, Beechman et al (2012) Beechman & Wood (2014), and el Esawey et al. (2015) the authors created visual analytic tools from data collected from thousands of cyclists in the great area of London. This work is also similar to services sold by Strava Metro (*Strava Metro Home*, n.d.). Both academia and private industry are trying to create tools to assist in identifying cyclists' behaviours in cities to aid city planners in providing better infrastructure for their citizens.

On the individual data visualization spectrum, Wood (2015) proposed a series of visualizations to inspect individual progress in group cycling events such as Randonneuring (individual long-distance cycling). Riders can visualize (and annotate) an overview of their entire ride, compare their progress against other riders (Figure 18) or even see when specific riders dropped from the events (time and location). Beck et al. (2016) proposed a matrix-based time-series visualization for sports (in their work, cycling) to support post-hoc performance comparison of athletes in cycling time-trial events. Kaplan et al. (2016) created a system to superimpose pedalling forces (the forces applied to the pedals by the rider in every stage of the pedalling cycle) into video recordings of stationary cycling sessions. Their work promoted a new approach to teaching and learning proper cycling forms to cyclists more intuitively and with a lower cognitive load than previous techniques.

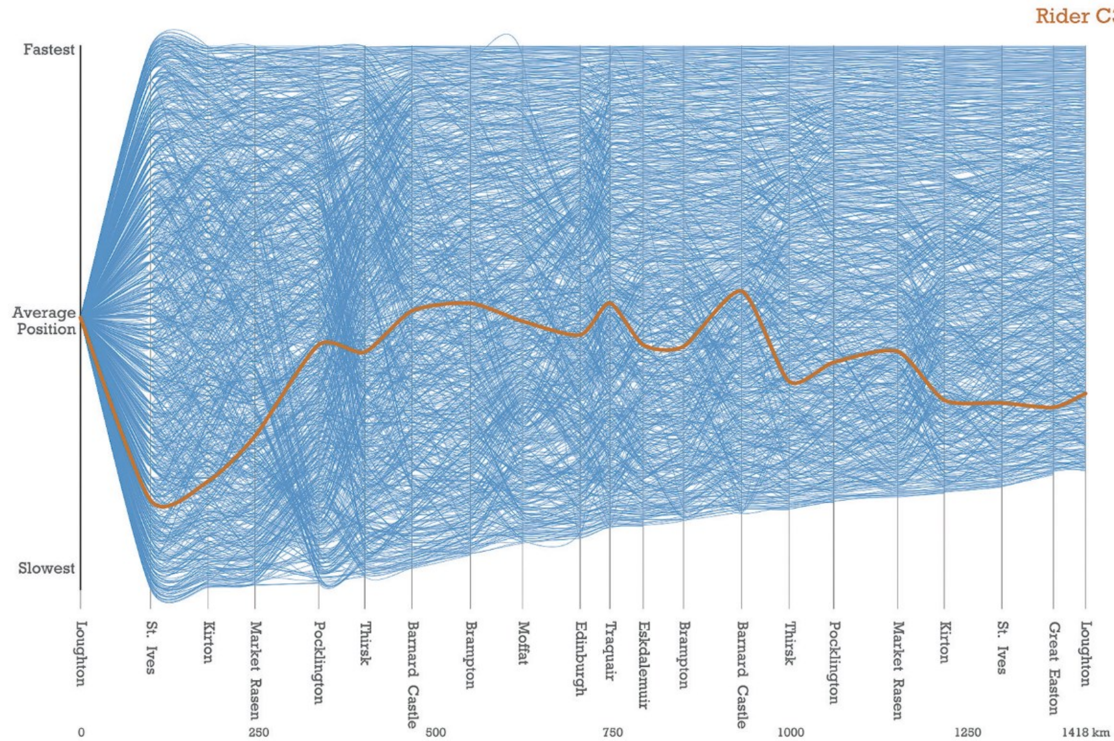

Figure 18 - Image from (Wood 2015) depicting the progress of a particular rider (C35) compared to others in a 1420km cycling event.

Outside academia, sport tracking websites such as Strava and RunKeeper (Track Your Run - ASICS Runkeeper Running Tracker App, n.d.) are used as a form of post-hoc data visualization and social media platform. Both websites allow athletes to share and compare activities with other riders among their networks. Visualization tools in such environments (Figure 19) consist mainly of a coordinated multiple view system combining time series analysis of cycling metrics with a 2D representation of the route and elevation.

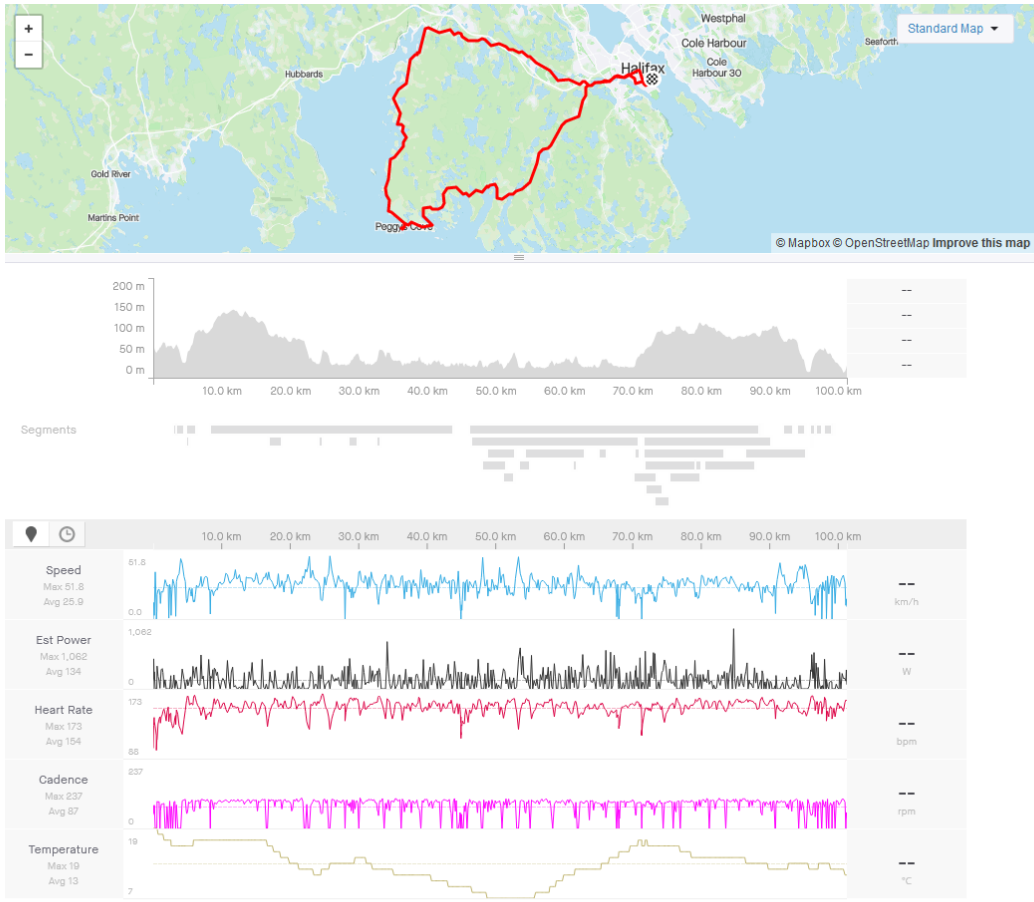

Figure 19 - Strava's Analysis visualization tool. Source: author's personal log

Although not with cycling, Stusak et al. (2014) created tangible visualizations artifacts (Activity Sculptures) that can represent a running activity. The artifacts are automatically generated and can be 3D printed directly from the data collected by sports watches (running and cycling share some amount of sports metrics). These sculptures are interesting as they break the traditional time series visualizations for sports data creating a less serious way of approaching the activities data.

5.2.1.2 Cycling Specific: During exercise

When riding, cyclists demonstrate a self-organizing behaviour, or as described by Waldron et al. (2011), a "swarm behaviour." This pack organization is highly evident in high-performance cycling, where riders train together daily. On the other hand, novices might take more time to self-organize or even rely on a more experienced rider to promote such conduct since in-ride communication protocols and awareness of the other riders' intentions and states are not yet understood.

There is an exciting design space in cycling (and maybe also in other team sports): how do we support sharing of personal metrics for each athlete in a group to promote distributed knowledge? Is it possible to combine individual data to create a distributed state of the group (i.e., Distributed Cognition), which will facilitate cyclists' self-organizing behaviour? Will achieving this group awareness be possible even without all riders having access to all state information? While these are interesting questions, most of the literature and commercial systems focus primarily on the data of oneself, ignoring the other riders and the surrounding context of the rider. Chapter 4 describes the current data sensing and visualization tools available for cyclists while they are exercising.

Previous research has mostly refrained from directly exploring visual tools for bicycle computers (probably because of the difficulty of deploying software packages to the devices); instead, researchers explored different methods to show and share data using purpose-built artifacts. Walmink et al. studied how riders can use the helmet as a visualization surface (Walmink, Chatham, et al., 2014; Walmink, Wilde, et al., 2014). In the latter, the authors explored how sharing the rider's heart rate data with other riders behind can enhance social experiences while exercising. Steltenphl & Bouwer (2013) moved the navigation aspect out of the bike computer displays by creating a belt with built-in vibration motors to improve navigation and, at the same time, increase safety, as riders can keep their eyes on the road while navigating from point A to B. An early work by Jones et al. (2007) demonstrated another application for helmet design but as an input device instead of an output. Riders used their helmets to create handsfree annotations such as a GPS-tagged road hazard, which could then be used to alert other riders in the vicinity that might be riding in the same road and direction. In similar work, Matviienko et al. (2018) proposed a series of bike and helmet augmentations such as vibrotactile motors and speakers to increase riding safety for young cyclists. In a controlled lab test (using a cycling simulator), they found that children exposed to the bicycle augmentations suffered no "accidents" while riding.

In Gesture Bike (Dancu, Vechev, Ayça Ünlüer, et al., 2015), the authors explored three different projection surfaces for navigation (HUD and on-the-pavement) and signalling (signal POD under the saddle and on-the-pavement) while riding at night. Like other studies involving cars and motorcycles, participants in their study preferred the HUD for navigation. They argued they could pay more attention to the road and traffic while being more straightforward to use than the pavement projection option. The NASA TLX score, used to measure subject workload, also showed a significant difference in performance, with HUD being higher.

Another noteworthy data visualization option for cyclists is using the bicycle's wheel as a type of persistence of vision display, as presented in AwareCycle (Kadomura et al., 2014). With AwareCycle (Figure 20), riders can share their cycling metrics with sports spectators (think of riders on a velodrome) or people on the street, without any special equipment on the viewers' side.

Figure 20 - AwareCycle (Kadomura et al., 2014) *presenting the rider's heart rate on the rear wheel*

5.2.2 Distributed Data Collection in Cycling and Environment Sensing

Bikenet (Eisenman et al., 2007, 2009) was one of the first works to collect bicycle and environmental sensor data from riders commuting in the city. It uses cellular networks and fixed access points to send sensor data opportunistically. Information is shared with the riders and with a remote server where others can interact with it.

Municio (2019) and de Brower (2018) propose a low-budget hardware device to collect sensor data on small or amateur cycling events. Their system does not share sensor data among riders (although it uses a mesh network where riders are the nodes). However, it focuses on collecting and sharing individual rider data with the race committee and spectators. Such work could drive augmented views of cycling data like those proposed by Lo (2019) on field sports. Cespedes (2019) proposed a prototype system to support the coordination of the travelling speed of riders in urban areas to promote safety. Their cruise control system supports ad-hoc join and aims to maintain a commuter group together to gain safety in numbers.

Jeon and Rajamani's (2018) targets sensing the rider's environment to promote safety. Their work explores the design and impacts of a rear-view car sensing system to provide collision warnings for riders and drivers at the same time. Smaldone (2011) presents a similar work detecting cars using video and audio from behind the rider to identify cars approaching in unsafe manners that could lead to a collision. Their work further extends the behind the rider sensing to sensing and storing road conditions ahead of the rider, such as potholes and unmaintained asphalt. On the same topic, Beecken et al. (2019) developed a system to classify road surfaces using a video stream from when riders are commuting. Regions of the city are mapped through crowdsourcing of multiple cyclists and uploaded to OpenStreetMap for the public to plan routes.

5.2.3 Heads-Up Displays

An interesting approach to data delivery in cycling could be the use of heads-up displays instead of traditional HDDs. Heads-up displays are used as a solution to provide information to the user while not blocking his or her field of view. Traditional approaches can use a small display in peripheral vision or a semi-transparent glass-like surface that has information projected to it. Commercial solutions range from everyday use such as the Google Glass and Focals 2.0 to sport-specific approaches such as the ReconJet, shown in Figure 21 (now defunct), Vuzix Smart Swim and the RideOn Ski Googles.

Figure 21 - ReconJet Cycling and Running smart eyewear. Source:

https://www.engadget.com/2015/07/17/recon-jet-review/

Authors in automotive and motorcycle UI have begun exploring the advantages and disadvantages of HUDs for riders and drivers. Liu and Wen define in (2004) HUDs and HDDs as: “A head-up display (HUD) is a display that provides information via digital graphics close to the driver’s natural line of sight. A head-down display (HDD), in contrast, is a display which provides information below the driver’s line of sight, usually on the dashboard or the middle console in a location which forces the driver to divert their visual attention away from the road”.

Most HUD papers published on the ACM DL are split evenly into three areas: HCI, Automotive, and Graphics and Technology (e.g., SIGGRAPH). While some sports examples, such as the previously discussed work by Dancu et al. (2015) exist, outdoor studies that explore data delivery and safety are more common with drivers and motorcycle riders. Automotive UI researchers have conducted studies with HUDs regarding their effectiveness and safety compared to Heads Down Displays. Road cyclists share common traits with other road riders; therefore, findings from such areas are highly relevant to this work.

Kenichitor Ito (2015) and Tetsuro Ogi (2015) explored the design and evaluation of a HUD for motorcycle riders. The authors state that motorcycle riders are different from car drivers as they constantly scan the road surface for deformities and hazards. Such deformities, although sometimes small, are enough to de-stabilize the riders' heads, thus making it difficult to read information from the traditional handlebar-mounted HDD. These findings are similar to a study by Trefzger (2018) where the authors concluded that commuter cyclists spend up to 45% of their time scanning the pavement with their eyes.

Smith et al. (2015) conducted similar work evaluating the driver's performance when conducting secondary tasks, such as reading textual information comparing HDDs vs HUDs. In a similar approach, Jose et al. (2016) compared HDDs vs HUDs vs HMDs in a simulated driving environment. Their results concurs with previous work, where participants preferred HUDs over HDDs, had lower navigation errors while using HUDs, and, notably, their participants reported an increased awareness of hazards and the environment when using HUDs. Hazard awareness is essential when driving a motor vehicle. However, it is even more important when riding a bicycle on the road: road dangers such as potholes, fallen tree branches, wildlife, and even car/truck parts have a higher probability of incurring damage and injury in riders than on drivers due to different physical characteristics of both transportation modalities. The reader must understand that road cycling injuries may lead to riders not being able to ride for the rest of the cycling season, thus invalidating all the training done up to the crash, or to even more permanent impairments.

Researchers have also studied the concept of adding see-through displays to motorcycle helmets, providing a true HUD within the rider's line of sight. Häuslschmid et al. (2018) investigated HUDs in motorcycles from a safety perspective. Motorcycles are interesting, maybe even more than cars, as they are closer to human-powered bicycles. The authors explored the perceived workload of their participants when using the Driver Activity Load Index (Pauzié, 2008) and found that their HUD approach significantly reduced the perceived interference (from the driving tasks), situational stress and visual demand when compared to an HDD.

More recently, Topliss et al. (2019) also found benefits for HUDs over HDDs while driving. The authors showed significant results regarding the driving dynamics of their participants. Participants exposed to the HUD had higher lane-keeping scores and higher Minimum Time to Collision. Their results show the benefits of HUDs regarding the safety of drivers, which can possibly be generalized to riders, but further research is needed.

It is important to note that HDDs (in cycling) and HUDs are very limited in the amount of information they can deliver to the user due to their small screen and resolution. However, it might be worth exploring the use of a combination of both display methodologies, one complementing the other, instead of the current commercial approaches, which promote the use of HUDs instead of HDD. Adding a second display layer to road cyclists, where information is contextually placed on each display depending on the rider's current task or position in the paceline, can have exciting perks, especially if the beneficial results from other fields are transferred into cycling. For example, Medenica (2011) showed that "AR HUDs have been shown to help keep drivers' eyes on the road with lower levels of mental workload, and consequently improved driving performance." Furthermore, drivers respond faster to dangerous situations when using head-up displays (Horrey & Wickens, 2004; Smith et al., 2015), and they seem to have a preference for HUDs over HDDs (Jose et al., 2016).

5.2.4 Context-Aware Computing

Interacting with heads-down displays such as bike computers when exercising has an inherent risk. Debris, wildlife, and even other riders can suddenly become an impact hazard if a rider is not paying attention to the road or does not have both hands on the handlebar because s/he was interacting with the data.

Bill Schilit originally keyed the concept of context-aware software (Dey, 1998; Schilit et al., 1994). The authors define context-aware computing as “the ability of an application to discover and react to changes in the environment they are situated in.” This initial definition, however, only focuses on one element of context: relative location between mobile and stationary devices and other elements inside a room or rooms. The same authors define context as where you are, who you are with, and what resources are nearby. However, there are still other untapped elements of what constitutes context, such as external actors, environmental conditions, and even some measures of the inner state of each individual interacting with an application.

Pascoe et al. (1999) define context-aware computing as a paradigm for mobile computing where applications must be able to sense and react according to external stimuli. Ideally, applications should adapt to external changes independently (Weiser & Mark, 1993) without requiring explicit user input.

Context, as mentioned, is much more intricate than *just* location plus the people around you. Different authors tried to define context's variables and transitions by aggregating different aspects of how humans interact with computers and with others around them. Anind Dey (1998) defines context as a combination of the user's emotional state, location, date and time, and objects and people in its environment. From Brown's (1996) simplistic definition of context: "any elements of the environment known by the computer", Dey (2001) further refined it as any information that can be used to characterize the situation of entities that are considered relevant to the interaction between a user and an application, including the user and application themselves. Lastly, Dourish (2004) defines context as an ad-hoc element of the interaction between two or more people; it is unique to a particular interaction and ceases to exist the moment the interaction stops or changes. Context arises from the interaction itself and is maintained by its unique details. Therefore, context-aware computing can be defined as systems that provide relevant information and services to users based on their current context.

As expressed by Dey (1999), context-aware applications, however interesting and probably useful, are difficult to build due to the intrinsic complexity of context. Challenges arise in sensing, classification, framework support, and context delivery. The work presented in this thesis assumes an external entity (such as a context server (Dustdar & Rosenberg, 2007)) responsible for detecting and providing contextual information to the proposed visualization system. The scope of inferring context from the rider's position and environment is beyond the scope of work.

Researchers are also exploring the usability and acceptance of context-aware systems. In e-graffiti (Burrell & Gay, 2002), the authors evaluated an application that allowed people to create and read virtual world graffiti in real-world locations (position-aware context). They observed that although participants initially found the idea appealing (probably due to a novelty effect), the application was not used considerably during a long-term study. One reason for the low usage might be that participants had to actively search for notes instead of being prompted by the application when notes were nearby. Modern location-aware approaches such as Google Assistant[12], Siri[13] or Yelp[14] combine other context variables with location to provide notification of nearby places that might be of interest to a given user. Burnell and Gray also observed that participants used the e-graffiti differently than the initial scope of the application and suggested that location-awareness was not something expected of a computer application at the time.

Fast-forward to the past couple of years, location-aware applications are much more common, from games such as Pokemon Go![15] to dating applications such as Tinder[16]. Context-aware mobile applications are not only position-aware but also use a plethora of variables to establish users' current context.

[12] https://assistant.google.com/
[13] https://www.apple.com/siri/
[14] https://www.yelp.com/
[15] https://www.pokemon.com/
[16] https://tinder.com/

Throughout time, researchers had access to a more significant set of sensors and algorithms to analyze the data provided by them. Context sensing is not restricted to location-sensing anymore, although it is still one of the main drivers of context. In InterruptMe (Pejovic & Musolesi, 2014), the authors evaluated how context can be used to find opportunistic moments to interrupt users with notifications from their mobile phones. They defined context as a combination of time, discrete location (work, home, public space), and accelerometry data from the user's smartphone. The authors found that although location combined with phone sensors provide a good starting point, the context for interruptions should also consider other aspects such as participants' current company, emotional state, and the subject (work, leisure, adds) of the notifications. Examples of other sources for context inference include speech, activity sensing and classification, nearby objects, gestures, and physiological state.

Examples of context-aware systems that go beyond location include the context-aware recommendation system for event-based social networks by Macedo et al. (2015), where the authors defined context as a combination of the user's current group memberships on social media, distance to any given event, and history of the user's previous preferences such as event description and number of confirmed RSVPs in order to recommend events to any particular user. In Just-For-Me (Adnan et al., 2016), the authors propose a context-aware music recommendation system by sourcing the user's location and historic music preferences, combined with global music ratings, e.g. top hits in Canada or NS, and music content as the driver for context. Context-aware systems are also gaining traction in tourism research: Braunhofer et al. (2015) define context as a combination of travel time to a specific POI, current weather, user's available time, and their POI visit history. The authors then used this contextual information to make suggestions of places that the users might be interested in visiting.

The examples above demonstrate that the definition of what constitutes the user context is intrinsic to each application, although some variables such as location are present in multiple examples. From what was discussed above, it is also possible to observe that digitally sensed context, as opposed to real-life context, is becoming denser as better sensors and algorithms, e.g., machine learning and inference, are emerging.

5.2.5 Road Cycling Summary

Yvonne Rogers defines context-aware computing as "[a system that] focuses on detecting, identifying and locating people's movements, routines or actions with a view to using this information to provide relevant information that may augment or assist a person or persons" (Rogers, 2006). Ideally, computer systems focused on delivering data to athletes should be aware of the environment and the current context (e.g., position in the line, cars, grade). They should display the required information with minimum interaction from the athlete and in a way that makes sense at any given moment (e.g., choosing between displaying information on the computer or a heads-up display based on the information type and the riders' position in the line).

For athletes, such as cyclists, computer systems already sense much of their environment with their body-worn and bicycle/shoe-mounted sensors (some bicycle pedals can even register if the rider is sitting or standing). These tools are ubiquitous to enthusiasts and professional sportspeople, but besides displaying the information from the sensors in a grid-like visualization, nothing more is done regarding the data. An information delivery system that can infer context from the sensors and the environment to identify the best way to display information and possibly reduce the number of times riders need to remove their hands from the handlebars to interact with the data would be invaluable in promoting riders' safety.

Interestingly, to the best of my knowledge, there are no examples of research and applications aiming to provide context-aware data visualizations for athletes during training sessions. The only exception seems to be Strava’s Live-Segments (*Live Segments – Strava Support*, n.d.) which provide in-exercise location-aware notifications of upcoming community segments[17].

In summary. while the amount of sensing and computational power in cycling has increased over time, very little has changed regarding data visualization and interaction. Devices still use the same grid-like data view (more discussion in Chapter 4) inherited from the first "analog computer" in the 80s. Interaction with modern computers is still based on either side buttons or directly touching the display, which may be a safety issue depending on the situation.

Modern devices are equipped with Bluetooth, WiFi and ANT+ radios. Nevertheless, there is no information sharing among riders on the same paceline, which creates information silos hindering a distributed and informed decision-making process. Lastly, manufacturers have tried, with little success, to completely replace HDDs with HUDs instead of pairing them and trying to leverage the unique advantages of each display technology.

[17] “[Strava] segments are portions of road or trail created by members where athletes can compare times.” (*Live Segments – Strava Support*, n.d.)

5.3 Methodology

I followed the Contextual Design (Beyer & Holtzblatt, 1997) methodology to model the work done by cyclists, as it provides a means to study individual and group behaviour from multiple perspectives.

In Contextual Inquiry (CI) (the first stage of Contextual Design), participants are observed while working in their natural environment. Researchers observe and query participants, forming a master-apprentice relationship, to learn as much as possible from participants as they engage in their regular activities.

5.3.1 Contextual Inquiry During the COVID-19 Pandemic

COVID-19 restrictions prohibited close interactions between individuals, which is expected in traditional CI. Group sports were also cancelled in most jurisdictions during this time. To address this issue, I followed an adapted version of CI, using data collected by the riders, which I used as a tool to promote recall of their rides during a remote contextual inquiry/interview. This approach is inspired by Beyer and Holtzblatt's guidelines for extremely long tasks:

> *"[…] do a work walkthrough, which is like an in-depth retrospective account. Set up an event in which customers bring in project documentation from all parts of the process and walk through the history of the project, week by week, meeting by meeting. Use the project artifacts to ground the inquiry" (Beyer & Holtzblatt, 1997).*

5.3.2 The Remote Contextual Inquiry

5.3.2.1 Participants

I recruited participants using social media and regional governance bodies in Canada (e.g., Bicycle NS). Participants needed at least two years of group ride experience and to ride at least 100km weekly during peak season. Participants were recruited in groups of 3 or 4 from the same riding pack. The packs could themselves be bigger than three persons.

5.3.2.2 Materials

Participants were asked to select two solo rides (rides they rode alone) and four group rides. The four group rides were the same for all participants from the same pack, allowing me to cross-reference shared experiences. Participants also shared pictures of their bicycles, artifacts they carry with them when riding, and one image for each page of their bike computer (Figure 22).

I conducted the Remote CI interviews on Microsoft Teams using voice and video. At the beginning of the first encounter, participants had the opportunity to ask any questions related to the study and consent form before giving verbal consent. I compensated participants fifteen dollars for the first encounter and ten for each of the following meetings. My modified contextual design process involved three contacts with participants: two individual interviews and one group interview with all members of the same riding pack.

Figure 22 - Pictures of four bicycle computers of participants in the Contextual Inquiry. Different participants had a different quantity of data fields selected.

The first interview was the contextual inquiry component, lasting an average of two hours for each participant. The protocol for this interview was as follows: first, participants had an opportunity to talk a little about their experience and motivations behind participating in the sport. Then, participants would explain in general terms their bike computer screens and data fields while explaining why they chose to have those specific data fields on each page, which page is used the most and when they use each specific page. Participants then explained to the researcher how they communicated with other riders while riding and demonstrated their gesture dictionary while explaining each gesture.

Once this last step was completed, the researcher would introduce the concept of contextual inquiry again and reinforce the master-apprentice relationship by stating that participants are the masters and the researcher is there to try to learn as much as possible. Participants would then go over each of their individual rides while the researcher tried to promote recall using the data available on Strava. The goal of using the solo rides was to provide a stress-free (explained to participants as training wheels) opportunity for participants to experience recall and storytelling without risking information loss in the group rides. Participants were informed that they did not have to go over their rides linearly; that is, they could start from the middle or the end or jump from one point in time to another if they wanted to. The researcher, however, would always start each ride-specific interview by asking for weather information and preparation steps before each ride (such as food packing and planning of route) to try to promote recall and immerse the participant in the context of the ride.

After the individual rides were completed, the researcher exposed the participants to the group ride data and a similar process to the individual rides was followed during the interview. Once the group rides were completed, the researcher would go over their notes and ask questions that they might have. This last step was also another opportunity for an open-ended discussion with the participants regarding their group riding habits.

I repeated the same process for each participant. This was not only an opportunity to cross-validate individual findings with others from the same pack, but also a chance to find issues in communication among the participants. For example, one participant stated that when they were leading a particular section of the road they felt as if the pace was fine for everyone in the group, while another rider from the same group said that on that same section they were working harder than they should be to keep up, and were not on active recovery.

After the first interview, I created the five Contextual Design's[18] work models for each rider: Culture, Sequence, Physical, Artifact, and Flow. I then invited participants for their second individual interview. In this interview, I explained the purpose of each work model and presented them back to the participant. I then reviewed the information in each work model with the participant to validate and refine our models. Participants used this opportunity to propose changes and add extra information, including examples from other riders that they might have remembered. This second interview lasted forty-five minutes.

With all models validated with their respective participants, I collected the work models and generated consolidated work models for each pack. I then brought the participants from the same group together for a final meeting. In this meeting, I explained the work models again and presented the consolidated models. Participants had the opportunity to validate the consolidated models as a group and discuss their findings. This step lasted forty-five minutes to one hour.

After all of the meetings were completed, I consolidated all group models into a final set of consolidated models (depicted in Appendices E – I) as a general representation of the work done by road cyclists riding as a group.

[18] The reader can refer to page 66 for an explanation of each model.

5.4 Results

I recruited twelve participants from three different cycling groups in two regions of the country. Members of the first group focused on long-distance cycling (Randonneuring); the second group focused on training for fitness and some race events; the last group was from a formally organized, coached cycling group. All participants rode primarily on asphalt surfaces. While participants had different backgrounds and years of experience with road cycling, they all mentioned practicing the sport because they enjoyed it and for health reasons, that is, to stay fit.

In the following subsections, I present insights from the Contextual Inquiry and Work Modelling activities. Participants are quoted using pseudonyms and are not aggregated by their cycling group to avoid reverse identification.

5.4.1 Roles and Responsibilities

One expects a rigid or formal role structure in a traditional office environment. For example, a developer is not expected to take the role of the accountant one day and custodial on the other. Interestingly, roles within road cyclists are fluid, albeit well-defined. I identified three "and a half" roles (the half part being a sub-role discussed further in this section): leading, middle-of-the-pack, and wheel (or last).

Cyclists are expected to know how to behave in each of these roles as they move between them constantly as part of the natural rotation strategy of road cycling. During a regular ride, an individual rider will be leading the pack multiple times. Unique tasks and responsibilities accompany each role. These necessary actions have an impact on the mental state of the riders. Participants mentioned being mentally focused and taxed when leading versus relaxed when at the wheel. Riders reported that the impact on their mental state comes from the responsibility that each role has associated with the pace and safety of the rest of the group.

> *"I do [worry about my mates behind when leading]. I always wish there was a better way of understanding how much they enjoy my speed, and I never asked. And if it's appropriate, if it's OK. I always wondered silently: Is it too slow? Is this too fast?"[Dave]*

> *"I wish there was a better way for me to understand how much they are enjoying my speed/pace. Is it too slow or fast?" [Charlie]*

Participants reported that the leading role had the most responsibility and mental effort. They mentioned focusing on upcoming debris and hazards, landmarks, wind direction, their computer numbers, and the grade[19] of the road. Wind direction was used to position the group on the shoulder (and sometimes on the entire traffic lane) to facilitate echeloning[20] when there was room to do so. Leaders were also responsible for adapting the pace when the grade changed, such as an upcoming hill.

Riders in the leading role must also monitor their effort using any data available such as power, speed, and body sensations. Participants mentioned monitoring their bodies closely, sometimes before looking at the numbers. Burning sensations, respiration rate, and other bodily indicators were all used as implicit measures of effort. The current leader is responsible for using this effort information to control the group's pace so that no one gets dropped from the paceline while everyone gets exercise stimuli. Information flows from

---

[19] the level and gradient of a roadway determined along the center line [Merriam-Webster]
[20] organizing the bicycles diagonally from one and another to deal with cross winds

the leader back to the leader without any external exertion information coming from the riders behind. I present the communication between riders in more detail in the Between Rider Communication Section below.

Furthermore, the current leader is also responsible for keeping the group safe. This includes ensuring hazards are signalled early in advance, positioning the group on a line that avoids bad asphalt, and ensuring that the group will only enter the traffic lane if no cars are coming from behind.

Unfortunately, not all elements of the tasks associated with the leading position are supported by the bike computer, making it hard for the leader to offload some of the responsibilities to it. Participants mentioned that it is difficult to find a pace that fits everyone when leading. They wished to know the previous leader's pace without having to use recall and, if possible, to monitor any gaps that might be forming behind or the physical state of their peers behind them. Observing the riders behind is difficult as it requires the leader to look back, thus not keeping an eye on the road ahead. This information on gaps is vital when dealing with stronger paces (to avoid dropping riders and creating a big gap between the group and the spent rider who got dropped), but it does not readily flow to the leader.

The second role I identified is the middle role. Riders who are not leading or at the wheel are part of this role. In the middle, I further identified a sub-role: the rider directly behind the leader has an added set of responsibilities. Participants defined the middle role as having a high mental effort, similar to when leading. Riders reported that one of the reasons for the mental effort is the added responsibility of not breaking the pack in the middle:

> *"When I was in the middle, I felt much stronger responsibility too. Keep my wheel closer to the person in front of me because now suddenly there is two people behind me depending on me keeping that draft, you know?" [Jake]*

Participants reported being attentive to the person's wheel ahead, trying to keep gaps small and consistent while avoiding overlapping wheels. Focus is also on relaying incoming hand gestures from the front to the back. Interestingly, participants reported not being too worried about cars while riding in the middle. These riders are not necessarily focused on their sports metrics or pace since they have very little control over the situation. This changes when the rider in the middle is just behind the current leader. The rider will start paying attention to the current pace of the group using the available information, such as power and speed, so they can be ready to maintain a similar pace once they become the new leader. Participant Mike summarized these extra tasks as steps to answer the following question *"How will my responsibilities change once I am the leader? Especially from a speed and power perspective"*. Because there is no sharing of cycling metrics among riders, it is up to the person immediately behind the current leader to manually keep track of the current pace, which they should maintain once the leader rotates back.

The third role is what the participants defined as the wheel or back. A rider is in this role when they are the last person in the pack. Participants mentioned being mentally relaxed as they have no one behind them. In this position, they allow themselves to wiggle from side to side and have the opportunity to grab food or other items from their jersey pockets. While keeping the gap small, they might drift back a little to stretch their legs off the saddle or make small sprints if the ride intensity is not high. Their only task is to keep track of any incoming vehicles and verbally inform the riders ahead if one is coming. Participants, however, mentioned that some of the questions they ask once they reach the wheel are currently not easy to answer:

> *"How was my pull? How was my pace? Was the pace appropriate for the group?" [Charlie]*

5.4.2 External Influences

Participants mentioned four key influencers: sports motivators, other riders from their group, road drivers, and provincial laws and infrastructure. Each contributes to a subset of the riders' actions, decision-making or feelings.

Participants mentioned cars as a cause for concern as they can cause dangerous accidents. The concern is not only for cars driving on the road but also for cars parked close to the shoulder of the road. The groups mentioned taking the road and moving into the traffic lane when cars are parked close to the shoulder to avoid being "doored." Although collisions with cars might seem unlikely for non-riders, one group was hit by a car during one recorded ride. The presence of cars also influences the available strategies when dealing with the leader rotation, echeloning, and road hazards, as described in the Dealing with the Environment and Group Safety subsection.

The most significant influencer, however, is the interaction with the other riders from the group. A rider either imposes their pace and road positioning on the other group riders or follows the leader's pace.

Participants also mentioned trying to be *"tough"* or *"with an endurance state of mind"* [Carol] in relation to their peers as if suffering during the practice of a sport is something to be expected, and you should not show it to others. Interestingly, this can lead to overly taxed riders who might not complete the ride or slow down the group for longer than necessary. Participants acknowledged this issue but were not necessarily open to addressing it:

> *"I tell people to ask to slow down if the pace is too strong for them. I, however, do not ask if I am the one suffering in the back" [Charlie]*

5.4.2.1 Sport-focused social media

In some capacity, all participants used specialized social media websites, such as Strava (*Strava | Run and Cycling Tracking on the Social Network for Athletes*, n.d.), Training Peaks (*TrainingPeaks | Reclaim Your Race Day*, n.d.) or Ride with GPS (*Ride with GPS | Bike Route Planner and Cycling Navigation App*, n.d.). Riders revisit rides and compare the cycling metrics of other riders. Participants used the platforms to plan future rides and download pre-determined group routes for the day.

Three participants mentioned using Strava to give positive feedback to other riders (using a feature called Kudos, similar to a Facebook Like) and to draw inspiration from what other riders were doing.

5.4.3 The Bike Computer

Participants mentioned relying primarily on the bike computer to support their work. While they carry more items, like food, water, smartphone, spare tubes, and tire levers, these items do not contribute directly to the work we observed but rather serve as supporting resources and tools for unexpected situations.

The bicycle computer was the primary information source for the study's riders. Computers are small displays that were either mounted in the steam of the bicycle or just ahead of the steam using a fly mount (Figure 23). Participants who had the fly mount mentioned that they preferred it as they could still maintain some awareness of the road while glancing at their data.

Figure 23 – Example of a bike computer using a fly mount. Source: https://www.garmin.com/en-CA/c/sports-fitness/cycling-bike-computers-bike-radar-power-meter-headlights/

All participants mentioned that although computers are beneficial, they are also a hazard as it forces them to look away from the road to read information. When they interact with the display (for example, to switch pages), they must remove one of their hands from the handlebars. The risk of accidents is more significant in group rides since riders are closer to each other:

> *"It is a little harder to keep track of the data in a group ride. I have to keep an eye on the wheel in front" [Dave]*
>
> *"I have to pay attention to the others" [Bob]*
>
> *"You can't constantly look at the computer, you have to keep an eye on the road and on your partners to avoid accidents." [Eve]*

All participants but one used a fixed bike computer; the remaining participant explained that they use a smartwatch instead because they participate in triathlons but is considering buying a fixed computer and running both the computer and the watch in parallel during these events.

Participants reported having at least two data screens on their computers: a primary screen (data page participants reported using most of the time), depicted in Figure 22, and one or more secondary data screens with more specialized information. The primary data page consisted of a grid-like layout rendering numeric information from sensors. Although each rider had their primary page customized, they all had similar layouts with at least speed, cadence, power (if equipped), time, and distance; the average speed field was present in half of the participant's primary display. While most of our participants had a power meter, two did not and used heart rate instead, with the field in their primary display. Both participants mentioned wanting a power meter but were waiting due to the costs of such a unit.

Even though secondary pages varied from participant to participant, the overall idea was to have one or more specialized pages to be used depending on the task or current role in the paceline. Four participants had a dedicated page when leading the group. This page contained less information than their primary display and focused on power and heart rate information. Other secondary page usage includes navigation (map-based), elevation profiles, interval training (self-guided or following a pre-defined plan) and commuting-specific fields.

5.4.4 Between Rider Communication

All participants mentioned some form of environmental impairment in their tasks. Wind noise was mentioned by all participants as a limiting factor to verbal communication, reducing the range of their voice from nearby riders only to no one, depending on its intensity.

"When there is a lot of wind, [the] vocal component gets lost" [Charlie]

The flow of information within the riders in each role can be divided into three parts: information from the rider's computer back to the rider, forward to the back of the paceline, and back to forward. All sensed sport metrics are recorded and (some) are displayed on the rider's computer. This information does not propagate to other riders. Information on rotation starts, debris and other road hazards are always from the front to back using hand gestures and sometimes verbal utterances such as "hole." Hand gestures, however, can sometimes be missed (like a UDP datagram being dropped).

Communication from the back – such as information about people dropping or incoming cars – is verbal, as no digital communication tool is used, and hand signals flow from front to back. This is affected by environmental constraints such as wind and cold, as described by participant Ana: *"Verbal communication is not 100% reliable. It is* [city of the participant]*, super windy, and it gets worse in colder seasons because of physical obstructions* [like a toque]".

Participants communicate mainly using hand gestures (information travels from the front to the back of the paceline). The gesture lexicon varies slightly on the stopping gesture from group to group (as observed from the stop gesture on the list below). However, most gestures were the same across all participants.

- *Stopping:*
    - Closed hand behind back
    - Open hand behind back
    - Open hand down on the side
    - Open hand lifted on the side (at head height)
- *Cracks on the road:* Move hand parallel to the bike
- *Loose Gravel:* Shaking open hand on the side
- *Rail Track:* Move open hand perpendicular to the bike
- *Hazards (Debris, Holes):* Pointing at the hazard
- *Move out of the Shoulder:* Sweeping Gesture from Right to Left (right-hand drive roads)
- *Start Rotation:* Elbow flick (participants mentioned that some riders are not obvious with this gesture)

Gestures are primarily used to identify upcoming hazards such as road debris, potholes, obstacles like parked cars and runners on the shoulders, and train tracks (which can "catch" the wheel of a road bicycle and cause an accident). The dictionary includes gestures to inform turns (although these were rarely mentioned in the interviews), when the leader is starting the rotation from front to back, and when a rider is going to get out of the saddle.

As mentioned, gestures only work from the front to the back of the line. Riders on the back cannot send effort information and warnings about incoming cars using hand gestures.

Furthermore, a gesture may fail to propagate backwards because a rider missed the signal (e.g., they are using the bike computer or grabbing their water bottle) or because the rider did not feel it was safe to remove one of their hands from the handlebar. Whatever the reason, when a rider fails to propagate the information backwards, riders in the back are at a disadvantage from a safety perspective.

> *"You're looking down. Maybe it's something on your bike computer at the time that the hand signal comes up and you missed the person's hand signal in front of you. Or perhaps the person in front gestures too late and you're pretty much on top of the hazard or just you know, just missing the hazard and it's too late to point it out so people don't have time" [Dave]*

Participants also mentioned that the intensity of a gesture encodes the severity or risk associated with what the gesture represents. For example, when dealing with potholes or debris of different sizes, participants might put more or less emphasis on the gesture: *"The bigger the hazard the more forceful I point." [Charlie]*.

5.4.5 Dealing with the Environment and Group Safety

In a more traditional work environment, such as a desk in an office space, a Contextual Inquiry Physical Model represents how the physical environment supports or gets in the way of the work. It shows how people organize their environment to support their work (adapted from Beyer and Holtzblat (Beyer & Holtzblatt, 1997)). In cycling, there is support from the physical environment (for example, having a paved road with dedicated bike lanes). However, riders cannot easily modify their work environment and are more often at the mercy of any challenges imposed by it (bad roads, crosswinds, rain, cold, or heat).

Where riders choose to ride on the road is influenced by elements such as the shoulder's width, wind, the presence of parked cars, and how busy the road is with drivers. Participants mentioned being more relaxed on roads with broader shoulders, allowing them to *"find the line"* with less debris through the ride by moving left and right within the shoulder. In contrast, riders mentioned feeling more vulnerable on narrow shoulders and noted that they would often ride at the shoulder line or even *"take the lane"* at moments when they thought they could *"get squished"* by vehicles.

Traffic density also influences how riders deal with different wind conditions. Wind can be a foe or an ally for riders depending on where it originates. If it is coming from behind, it effectively pushes the riders, helping them ride faster or with less effort. When it comes directly ahead or at an angle, it is working against the cyclists. With diagonal winds, participants mentioned riding in an Echelon (Figure 24) formation to facilitate the drafting and reduce the impact of the wind on tired riders. However, this technique requires a decent amount of road real-state since riders must be diagonally from each other. If the road is busy with cars or does not have a reasonably wide shoulder, riders cannot echelon to save energy due to safety concerns.

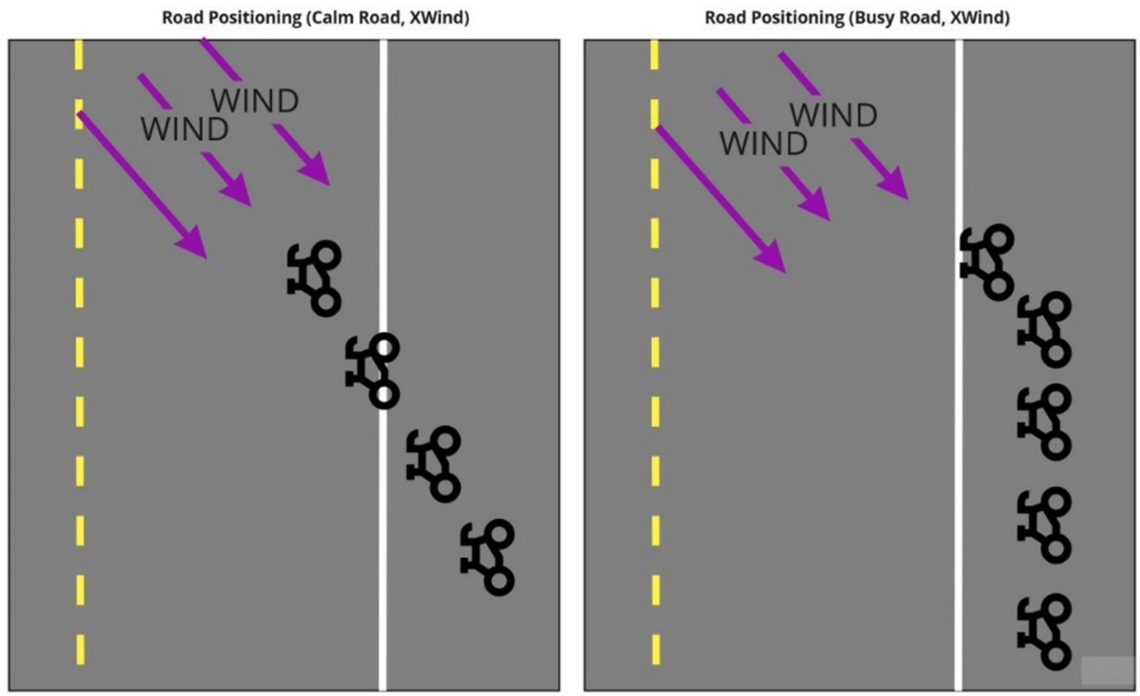

Figure 24 - Example of the echelon technique. On quiet roads, riders can use the entire road to deal with crosswinds. On busier roads, only the second rider can benefit from the configuration. If groups are large, riders will ride (and echelon) two-abreast when there is enough room on the road to support this configuration.

The presence of cars influences how riders deal with their rotation strategies (to share "the load" of moving the group ahead) and with road hazards. When the current leader identifies a hazard or debris in the road, the rider will try to move around to avoid hitting it. If no cars are close, the group moves into the traffic lane until they are clear to return to the shoulder. If there are cars in the traffic lane, riders will not be able to take the traffic lane and will have to slow down and wait for the cars to pass or, in a worst-case scenario, slowly hit the debris in front of them. The presence of cars in the lane will also block riders from rotating. As shown in Figure 25 right, riders leverage the traffic lane to rotate the leader. The leader moves into the traffic lane and waits until all riders go by their right side before moving back into the shoulder. If cars are present (or long hills and turns), the current leader will typically stay at the front of the group for longer than they should be, leading to further exhaustion and thus requiring a longer recovery time in the back.

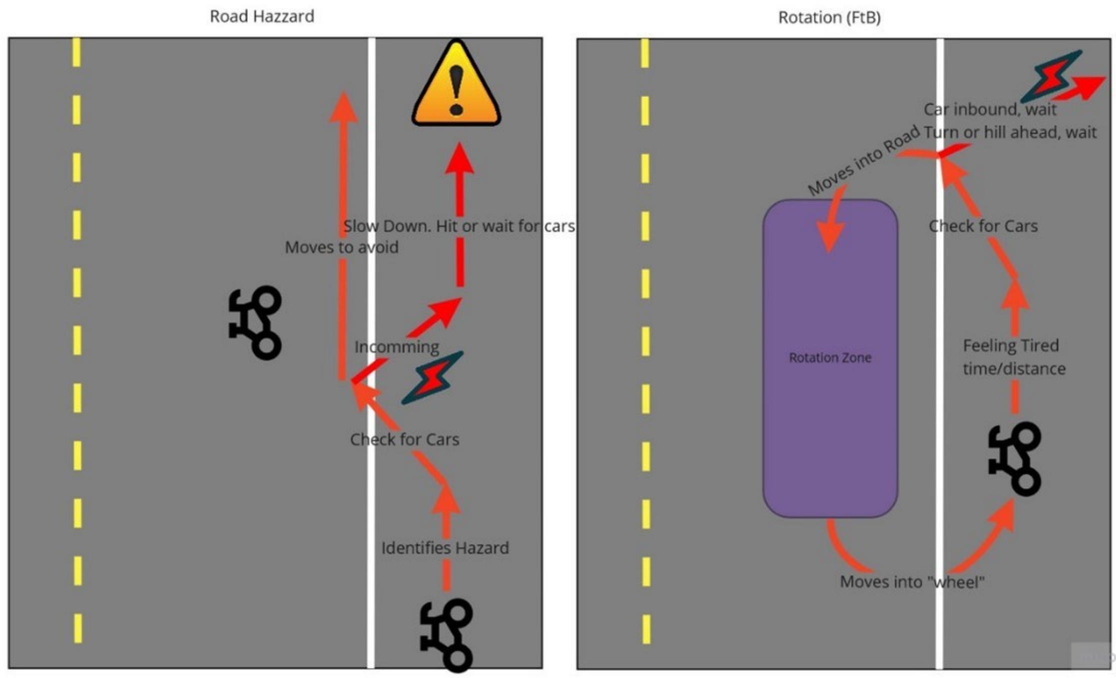

Figure 25 - Illustration of riders' strategies to deal with hazards on the shoulder (left) and Front-to-Back rotation (right). Arrows in red accompanied by a lightning bolt indicate that something blocked their intent, e.g., cars coming from the back.

5.4.6 Pacing Coordination

I observed two main group intents or goals regarding pacing from the interviews, which will be discussed in this section. They are:

- Rotating to "share the load" of leading and dealing with windage.
- Appropriately pacing the group to avoid overtaxed riders and drops

5.4.6.1 Rotation Strategies

Participants mention two common strategies to rotate the leader at the front of the paceline: back-to-front (BtF) and front-to-back (FtB). Although both approaches have the same end goal, they differ on who controls the situation, the current leader (FtB) or another rider in the pack (BtF).

Riders mentioned using BtF very rarely as it *"leads to less smooth rides"* [Charlie] and that "*it can be dangerous"* [Carol]. Group 2 mentioned using it when they started riding together for the first few times but have switched to FtB for the reasons mentioned above. Groups 1 and 3 mentioned seeing rare moments where a rider opts to take control of the rotation and start a BtF rotation. Another concern regarding BtF is that the leader has no control of the situation since the rotation begins behind them.

The overall steps for a back-to-front rotation are depicted in Table 6. It is interesting to observe the first pain point in red: because the leader is not the rider who starts the rotation, the leader is at the mercy of the riders behind. They had to wait until someone from the back noticed a drop in speed (indicating that the leader was probably tired), or someone decided to take over leading the group for another reason. Secondly, unlike an FtB rotation, the rider who wants to use this strategy must increase the pace to go faster than the leader to move into the front, thus expending more energy. Lastly, finding the correct pace when assuming the new leadership position is not straightforward since the new leader has now to deal with wind and the previous exertion of going from the back (or middle) to the front.

Table 6 - Sequence Diagram of the Back-to-Front rotation. Text in red accompanied by an exclamation mark illustrates issues identified by the riders or the researchers.

INTENT: BtF Rotation
Observe leader's speed
↓
Compare it with current implicit effort level
↓
If can do a better pacing take over; otherwise, wait until leader slows down
! No direct feedback on how tired the leader is.
↓
Check for cars, hills, curves
↓
Moves into the traffic lane with the entire paceline.
↓
! Picks up the pace as to pass the current leader
↓
Moves back into the shoulder after passing the leader
↓
Start pulling the pace line
↓
Maintain previous rider pace (speed)
! HR is not a good indicator of pace anymore since there is no more draft
! Sometimes the new leader unintentionally increases the pace
↓
Keep pulling until rider from behind takes over
! No information on when this is going to happen

In front-to-back rotations, the current leader controls when to start rotating to bring a new rider to the front. Participants mentioned two main sub-strategies: time/distance-based or exhaustion-based. Both are similar on most steps and differ only when to start rotating. Participant Eve mentioned: *"Generally what I tried to think about was how long the leader was pulling for. I tried to imitate that to be honest. And then in my head I think OK, you know, roughly they pulled for that long so it's kind of my turn to pull this amount"* as their strategy to start leading.

Both strategies are depicted in Table 7. The participants stated that rotating based on time or distance is the most preferred way in their groups. Interestingly, there is no computer-aided support for it, and it is the rider's responsibility to keep track of their pulling time or distance while trying to remember what the previous leader did. Participants mentioned using 5 min or 5 km as a starting point and would change the value depending on heavy traffic (longer pulls), hilly course (shorter pulls) or wind conditions (longer or shorter depending on direction) without explicit communication.

In both rotation strategies, riders will not begin a rotation if there are cars behind or if the terrain is not ideal, e.g. hills, curves, or bad roads.

Table 7 - Sequence Diagram of the Front-to-Back rotation. Text in red accompanied by an exclamation mark illustrates issues identified by the riders or the researchers.

INTENT: FtB Rotation
Leads until time/distance is up or is getting too tired to continue. Tiredness is implicit (e.g., breathing, burning legs) ! No support from computer to inform if pulling for enough time as compared to others
↓
Check for cars, hills, curves Trusts rear-view radar if equipped ! Will not initiate rotation if there are cars behind, hills and turns ahead
↓
Signal rotation with arms
↓
Moves into the traffic lane
↓
Keep soft-pedalling while the group advances on the shoulder
↓
Moves back into the shoulder behind the last rider (becomes the new wheel)

5.4.6.2 Pacing the group

Ensuring the pace is correct for the group is one of the leader's primary tasks. A good pace means that the ride intensity is such that no rider is overburdened (except for the leader in special situations such as a segment attempt). One participant defined the optimal pace as: *"So, in the end, you cannot really try to optimize the pace. You are just making sure you are within a threshold, that you are not killing anybody. Even if it costs not being at the optimal pace. I am just trying to keep the group together"*.

Two common intents exist when pacing (Table 8 and Table 9). Riders can lead without another cyclist dropping, or they can have a rider dropping from the line. Both variations start precisely the same: the leader starts leading and has to find a way to infer if the pace is good for everyone else in the group. Leaders will use whatever tools are available, such as speed, grade, heart rate, power, and personal knowledge of the other rider's strengths.

The difference occurs when a rider is forced to drop from the group because they cannot keep up anymore. A gap will start forming as the rider slowly moves away from the main pack once the rider is exhausted. Someone from the pack will inform the leader of this situation, and the pace will slow down so the rider can catch up with the group again.

> *"She was 100, 150 meters behind. Then I slowed down to wait for her to catch up. She only said it was too fast half the way through" [James]*
>
> *"[Researcher] The moments where you lost the draft while you are the wheel, you basically had no way of communicating and asking them to slow down? – [Doug]I did not; I just looked to get back on the wheel."*

Table 8 - Sequence Diagram of the pacing strategy used by the riders when they are leading. Text in red accompanied by an exclamation mark illustrates issues identified by the riders or the researchers.

INTENT: Pacing
Start leading the paceline ↓ Infer *implicit* exertion level of others ! *"Maybe I am pushing too hard/slow?"* ↓ Monitor self-perceived effort plus Power/HR ↓ ! Extrapolate own exertion levels to the rest of the group There is no direct feedback from the group. Previous riding experience helps ↓ Try to observe (by looking back) the rest of the group ! It takes time to see if it is too strong. Can lead to people dropping or depleting their reserves

Table 9 - Sequence Diagram depicting the steps when someone drops from the group because the pace was too strong. Text in red accompanied by an exclamation mark illustrates issues identified by the riders or the researchers.

INTENT: Pacing - too strong
Start leading the paceline ↓ Infer *implicit* exertion level of others ! *"Maybe I am pushing too hard/slow?"* ↓ Monitor self-perceived effort plus Power/HR ↓ ! Gap open, rider drops ↓ Someone from the back informs that there is a gap forming or that someone has dropped ! Usually, not the person who dropped

As seen by the pacing intents, inferring the other riders' effort and exhaustion state is an essential aspect of riding as a group. Participants stated that before thinking about the others in the group, they would first pay attention to their bodies and numbers and then extrapolate that to the other riders. Participants also mentioned that riding with their group members helps create a mental model of everyone else's limits and comfort zones: *"I definitely think that the more you ride with people that easier it gets to go out with them. It helps to know how tired they are and how much you can push. It is not only effort too. For example, with experience, you can learn how much in the middle of the road you can ride without making the other person feel uncomfortable. So, I guess this component that just take time " [Mike]*. Table 10 depicts the intent of inferring other riders' state while pacing.

Table 10 - Sequence Diagram illustrating the steps taken by the leader to gauge the effort state of the other riders. Text in red accompanied by an exclamation mark illustrates issues identified by the riders or the researchers.

INTENT: Inference of Peer State
Starts riding at a constant pace
↓
Pays attention to breathing, legs and other implicit information
↓
Monitor self-perceived effort plus Power/HR
↓
! Extrapolates own effort to others
↓
Looks back and observe other riders
↓
! Sees that a rider is frequently losing the wheel
↓
! Before losing the wheel, riders tend to "grab it with their lives", leading to unnecessary exertion!
Riders are more reactive (slow after they are exhausted) than proactive (slow down before being exhausted)
↓
! No communication from the tired rider to the leader until max exertion is reached
↓
Reduces pace slightly to accommodate
! No direct feedback of how tired other riders are or if they have rested and can go a bit stronger

5.4.6.3 The Implicit Pace

The overall pace (how strong or easy riders are pushing) was not strictly defined before a ride by our participants, and it seems to be implicitly defined by the current configuration of riders in the group at any given ride:

> *"Now we didn't discuss what the pace would be [...] It's just sort of, just go out and see how people are doing; if people are not doing, you know,*

having a great fast day, it's generally accepted that we'll ride slower" [Ana]

I observed this pace adaptation phenomenon firsthand in one of the groups (the overview of the interaction is depicted in Figure 26). Riders from one group were riding for four hours at a strong but comfortable training pace when they saw another rider from their group (who was not participating on that ride) ahead on the road, going slower than them. After a couple of minutes, they caught up with the rider (depicted in purple in Figure 26), thus changing the configuration of the members of the group. This new configuration led to an implicit change in pace (stronger), taxing some of the already tired riders in the group. These riders had to drop from the group as they could not keep up with the pace, creating a new pace configuration in both groups (the rider in purple remained in the faster group) and increasing their separation over time. After a certain distance, the rider in purple took a turn, changed roads, and left the original group. The riders who had dropped from the main group were able to catch up and were back at the pace they held before encountering their peer on the road.

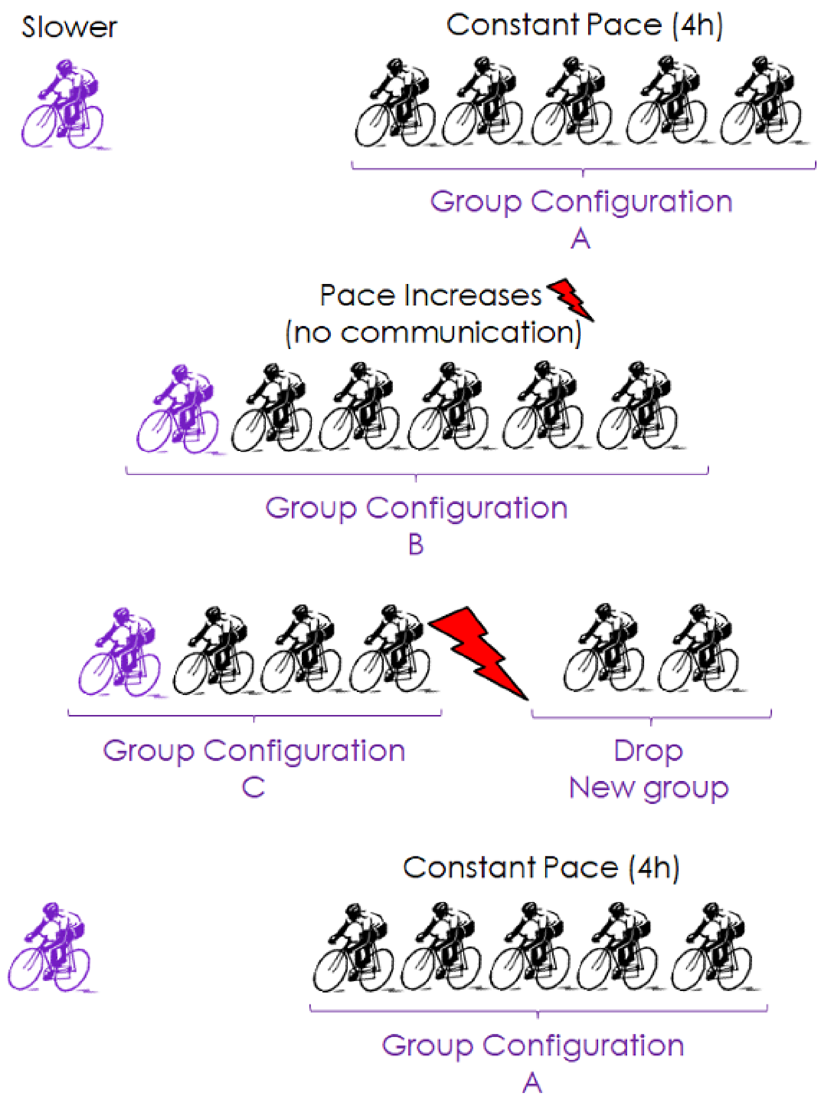

Figure 26 - Illustration of the implicit dynamic change of pace when the group configuration changes. Adding another rider to the group will impact the current pace.

5.5 Discussion

Results from the study show that Road Cyclists have needs related to collaborative actions, road safety (on a different scale than in-city riders), effort management, and sports metrics sharing. While the HCI community has explored tools such as smartphone applications (Savino et al., 2020), smart helmets (Matviienko et al., 2018) and augmented handlebars (Woźniak et al., 2020) with commuters and tourists, there is a lack of work that focuses on supporting the needs of cycling athletes. Although some of the needs are similar (e.g., road safety), others like group communication support, live effort sharing, and pace coordination are unique to the group cycling. Road cyclists have more restricted access to technology and due to higher travel speeds and close-quarter riding, interaction time with devices is shorter than commuters and tourists.

One core issue our participants elicited is how to communicate with the group correctly, promptly, and free of interference. While one might suggest, "why not use a radio?" radios are non-existent outside some specific pro-tour events. Even in professional events, radios are not used as two-way communication between riders in the same team. Radio channels exist as a one-way (simplex) tool between the team cars and the riders. Radios are not allowed in the Olympics and are non-existent in amateur events and training. None of our participants mentioned using headphones while riding with the group as it hinders their awareness of the road and the group. Only one participant stated they use headphones while riding alone but of the bone conducting type.

While some of the riders' challenges such as hand communication, flagging of debris and hazards could possibly be addressed by adding more hardware (e.g. smart gloves to detect gestures) and designing interaction layers on top of it, the bike computer has untapped potential to support the needs of the group plus it is ubiquitous. Although multiple issues were identified in the inquiry, I chose to focus the rest of my work on sharing exertion and maintaining the group together as it was identified as a key pain point by our participants. I believe that overtaxed riders and over/under pacing constitute a significant issue that can be addressed by exploring new interactive venues with the hardware already present with the athletes.

5.5.1 Pacing support and physical effort

From the results, group pacing is one of the tenets of group riding. Participants directly mentioned their strategies and issues in the Pacing Coordination Section and we can see how the lack of a good communication channel and environmental factors can get in the way.

On this topic, riders mentioned sharing the load to help manage the effort throughout their rides. From a resource perspective, we can imagine riders having a shared battery that they must maintain while exercising, and every rider contributes to it with their own batteries[21]. Draining the battery too fast or not letting it recharge a bit while riding in active recovery will inevitably get a cyclist overly taxed, creating the conditions for people to drop the paceline. Their primary method to control how much of the battery they are draining is for the current leader, using very little information from the other riders, to adjust the overall pace, thus managing the group effort.

Participants mentioned paying attention to information from their bike sensors (explicit data) and their bodies (implicit data). The combination of these two sources, each pooling from multiple sensors, can be used to gauge effort. From a systems perspective, it seems straightforward that sharing these effort metrics with other riders could potentially address

[21] The concept of a body battery is an analogy to the Work Capacity model defined by Skiba (2012)

the issue of overtaxing riders in the group. However, there are design implications on how to share such data and what format it should have.

From an individual perspective, explicit data is clear to read and make decisions based on it, e.g., if my power is increasing to a power zone higher than what I should be, ease up. Explicit data is also easy to share since the information already exists in the digital space. An exciting research challenge is assessing, quantifying, and sharing implicit data (like burning legs, laboured breathing, and tiredness from previous days). While one could only use power as a metric to pace a group of road riders, using body sensations might further assist riders in fine-tuning each ride based on more than *"just the numbers."*

Design should focus on how to represent this measure of effort to share with the other riders in the group. Walminnk's (2014) concept of sharing HR between two riders using a helmet display is an excellent start to promote socialization between the riders. However, it is unclear how it would scale to large groups. Also, absolute heart rate values carry little information because it is affected by ambient temperature and humidity, it is a *slow* metric that needs time to reflect changes in effort, and different riders have different maximum heart rates and ventilatory thresholds. Star's Boundary Object concept (Star & Griesemer, 1989) or the sports-related Extraction Objects (Kolovson et al., 2020) focus on supporting communication and collaboration between participants with different backgrounds or from different levels of an organization structure. However, the organization and roles for group cyclists while riding are flat. That is, all riders are the same and share formal roles with well-defined responsibilities. The responsibilities and roles change as riders rotate, and the information needed by each person changes at the same time. The social pacing rules between the riders are already in place. However, the digital tools to bridge the gap between the self and the group ("what I think of the group effort" vs. "I know the current group effort") are yet to be designed and evaluated.

From a group perspective, the effort is a combination of everyone's effort since everyone contributes to the group battery. As a result, interaction designers should consider who owns the effort metrics and who should have access to it. From one perspective, this belongs to each individual rider and is, therefore, private data; from another perspective, the data belongs to the group since everyone is riding together. Sharing effort becomes a fascinating design space problem: should we share an individual's metrics with everyone else, creating potentially uncomfortable situations such as exposing riders who are under-performing to their teammates? Or should we aggregate individual data into an abstract representation of group effort, removing the personal information?

From a privacy-preserving perspective, one could argue that not sharing individual information is the correct choice. However, ramifications of this decision go beyond privacy concerns only; therefore, research that considers both approaches (or a hybrid model) is required. For example, from a team-building exercise, knowing your peers' state might help you create a more refined mental model of everyone's capabilities and current fitness level. It can also help identify riders that might be more tired than others and leave them for longer periods closer to the wheel to benefit more from active recovery and drafting. The same could be true for stronger riders (or just less tired than others); they could spend more time leading the group, thus saving the other riders and contributing to a more balanced ride.

Although the above examples make a case for making individual metrics available, we should not dismiss the idea of sharing a group-wise metric that hides the individual data. While long-term group riders have a sense (and sometimes can even access historical data from previous rides) of their peers' capabilities, this is not true for more casual group rides where the group is formed ad-hoc on each ride. Forcing riders who are unfamiliar with each other to share their current effort while riding with strangers could make riders feel insecure and even lose interest in riding with others. However, an anonymized metric could still be used to manage the overall pace of the group to avoid taxing weaker riders while not singling them out. This metric could lead to stronger rides without drops, thus boosting morale (especially for weaker or more insecure riders) and possibly paving the road for creating a long-standing group of riders that will, over time, become familiar with everyone's physical abilities and goals.

The choice of privacy becomes situational or context-specific (Nissenbaum, 2004) to the current group's dynamics. It is defined by the cycling group's local constraints, such as their ride goals and friendship (or familiarity) with the other riders. Digital systems should support multiple ways of sharing the same (or similar) effort information based on the group dynamics of each ride.

I propose that private and anonymous options be available to riders on their cycling computers. The system could use power information and/or ratings of perceived exertion to drive the effort metrics. If we consider this group effort metric to be accurate, then other data becomes less necessary, simplifying the bike computer's display.

5.5.2 Too much data, not enough time

While modern bicycles can sense many sports metrics, communicating such information among the riders of the same group safely and sensibly warrants continuous exploration. Riders are paying attention to the road and their peers; therefore, it is crucial to consider how to deliver the information. Using a screen that requires more than a few seconds to read the important information might be a safety issue. Moving data to multiple pages can also lead to safety issues as riders will need to remove one hand from their handlebars to interact with the display.

> *"It is a little harder to keep track of the data in a group ride. I have to keep an eye on the wheel in front" [Ana]*
>
> *"You can't constantly look at the computer, you have to keep an eye on the road and on your partners to avoid accidents."[Dave]*

Role-based displays could potentially reduce the number of data fields rendered, thus making it safer for the rider to interact with the data. Managing information based on roles makes sense, considering that different roles have different tasks and information requirements. However, this concept creates another sensing challenge: identifying the role of the rider as it dynamically changes over time, which is somewhat unique to road cycling compared to other group sports. If we can sense roles, it is possible to offload some of the cognitive efforts to the computer. Information such as the pace and time of the last pull can be rendered automatically to the second rider when they assume the leading position. The leader can receive gap information (as in new gaps from a possible drop) if it occurs and be advised when their time or distance is over, and riders on the wheel can compare their pull with previous pulls.

Cycling computers could leverage the unique tasks and questions the riders have on the three and a half roles identified during the contextual inquiry. The information rendered should adapt itself to hide parts that are unnecessary, for example, masking group effort information for all roles but the leader since they are the only rider that can control the pace. Modern ad-hoc mesh networking and GPS positioning could possibly infer the position and role of each rider in the paceline locally. If heavier computation is required, 5G networks could support running role-detection algorithms in the cloud since riders don't necessarily need a high refresh rate of which information is rendered as their roles don't change every second but every couple of minutes.

5.5.3 A Design Exercise on Sharing

As a discussion exercise leading to the last study in my thesis, I created two bike computer visualizations to support sharing effort metrics with other riders (Figure 27 and Figure 28). The screens in Figure 27 share individual riders' agreement with the pace; riders can be OK with pace (= sign), asking for the pace to be stronger (up arrow) or slower (down arrow). The current leader still maintains the agency to decide to change the pace or not.

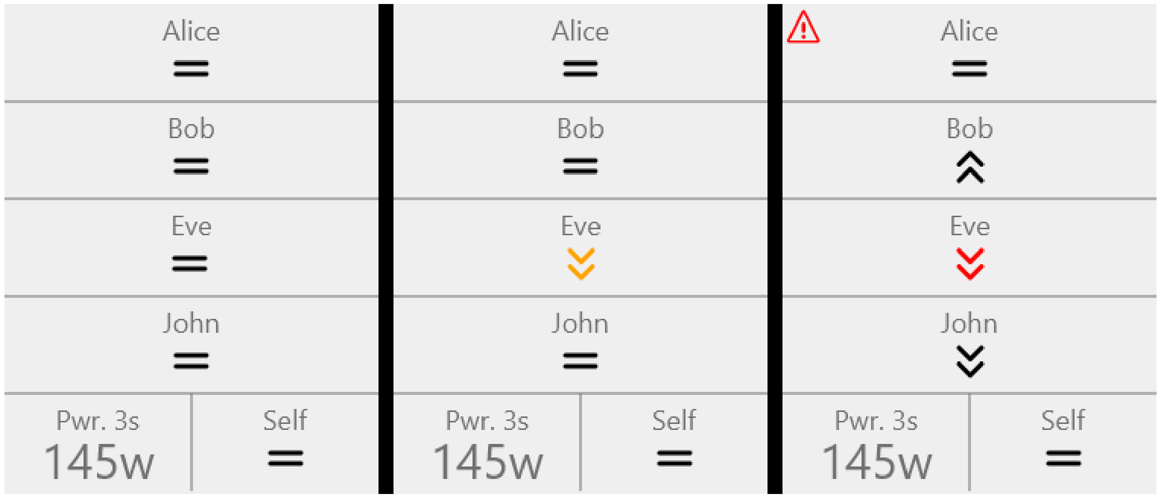

Figure 27 - Bike computer screen sharing individualized rider effort information in a traditional grid-like configuration. All riders have access to the same information. = sign indicates that riders are OK with the current pace; up arrow indicates that one rider wants the pace to be stronger; down arrow indicates that a rider is asking for the pace to be slower. The colour of the down arrow changes, over time, from black to red to indicate for how long a specific rider is pushing too hard.

Figure 28 shows the first iteration of the Pace Guide. It aggregates the current effort of every rider and combines it into a single anonymized representation deployed in groups where riders might feel uncomfortable explicitly sharing their effort. Only the green circle is lit if most riders are comfortable with the pace. Once one or multiple riders feel that the pace is either too strong or easy, the lights surrounding the green light will be lit accordingly. If at least one rider is struggling to keep up, all three top circles are lit (green, yellow, red) independently of the state of the other riders. The pace guide metric does not need much screen space since it condenses everyone's information into "one field." If riders want to have their traditional grid-like information on their computers, it is possible to overlay the data into their screens, as shown in Figure 29.

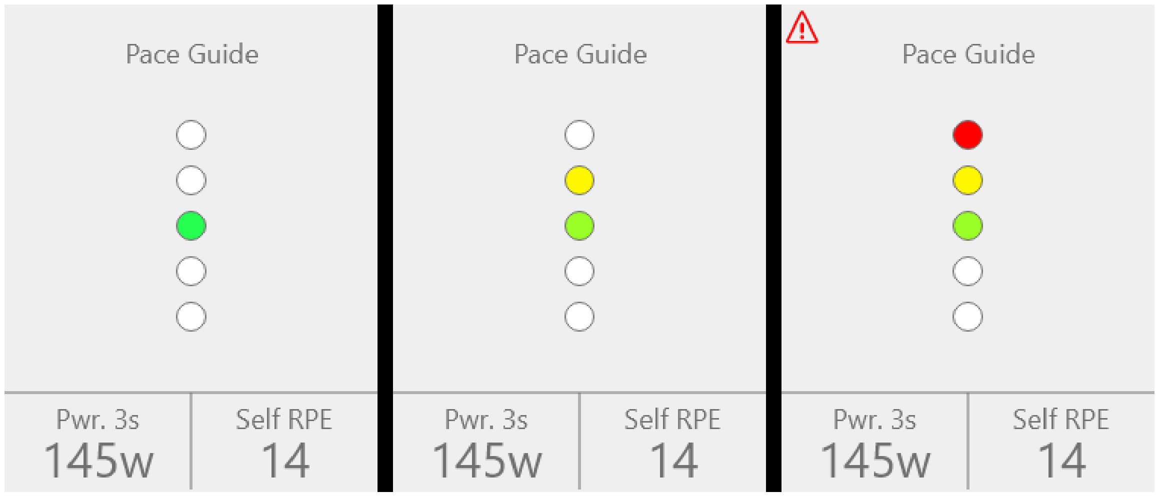

Figure 28 - Design of the Pace Guide metric. It renders all the group effort information into a single visualization thus preserving anonymity.

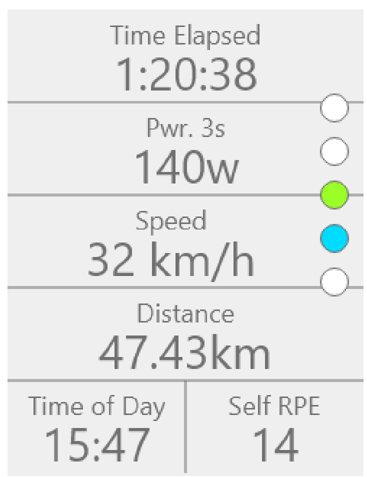

Figure 29 - Due to its small screen footprint, it is possible to overlay the Pace Guide metric on existing grid-like data pages without compromising any information.

5.6 Limitations

My study sample consisted only of Canadian cyclists. Although road cycling rules and techniques are the same around the globe, riders from different countries might have different cultural influences, potentially leading to slightly different models, e.g., cyclists in left-hand driving countries rotate in the other direction. All of our insights came from the ability of riders to recall their experience in the remote contextual inquiry. Even though we reached saturation with our three groups, there is the potential to uncover other insights with an in-loco contextual inquiry.

CHAPTER 6 IMPACTS ON SHARING RATINGS OF PERCEIVED EXERTION IN GROUP ROAD CYCLING

In the previous study, I identified challenges group road cyclists face when riding together. Communication between riders (duplex and multicast communication) seemed to be one of the highest points of attrition — more specifically, communicating how much exertion riders are putting out and how to manage strategies to maximize performance based on each rider's efforts.

My goal in the follow-up study presented in this chapter (ethics approval letters in Appendices K and L) was to explore how we can support the sharing of exertion information among riders in a cycling group and the impacts of doing so. Through this study, I sought to answer the research questions below. Addressing these questions could lead to higher coordination of the collocated riders and as consequence a higher performance output of the group thus leading to faster riders.

RQ1. What is the impact of the proposed effort-sharing techniques on participants' overall effort and energy expenditure compared to the baseline (how riders currently work together, with only individual tools)?

RQ2. Will participants be able to estimate the overall group effort better using the sharing techniques when compared to the baseline?

RQ3. What are the impacts of the sharing techniques on rider drops and overall time to recover from drops?

I explored the research questions through two interventions RPEView and Paceguide (Figure 33). Both stem from the proposed bicycle computer designs developed in the previous Chapter. I evaluated these interventions against a baseline condition, representing the current state of the art in commercial bicycle computers, in a controlled study using the virtual platform Zwift.

Results from the study show an increase in the power output of riders in the groups and a decrease in the numbers of riders being dropped from the group during the rides. Participants also mentioned being more aware of their peers exertion levels using both proposed interventions.

6.1 Paceguide and RPEView Exertion Sharing Tools

In the previous Chapter, I presented the work modeling created using Contextual Design. As per Dourish (2006) "Contextual Inquiry provides designers with a series of tools and techniques for understanding social settings and organizing their observations to derive models for design". I used my notes and models from the process to identify issues related to working together in a road cycling peloton. Chapter 5 concludes with a series of possible data pages for cycling computers (Figure 30) that focus on sharing exertion values (namely using RPE as the data metric) among riders in the group.

These data pages were created using my knowledge of the field combined with the discussion I had with participants of the Contextual Inquiry presented in the previous Chapter. They focus on providing the information in a single page (to avoid page scrolling) and to support at-glance visualization minimizing the time riders need to look away from the road.

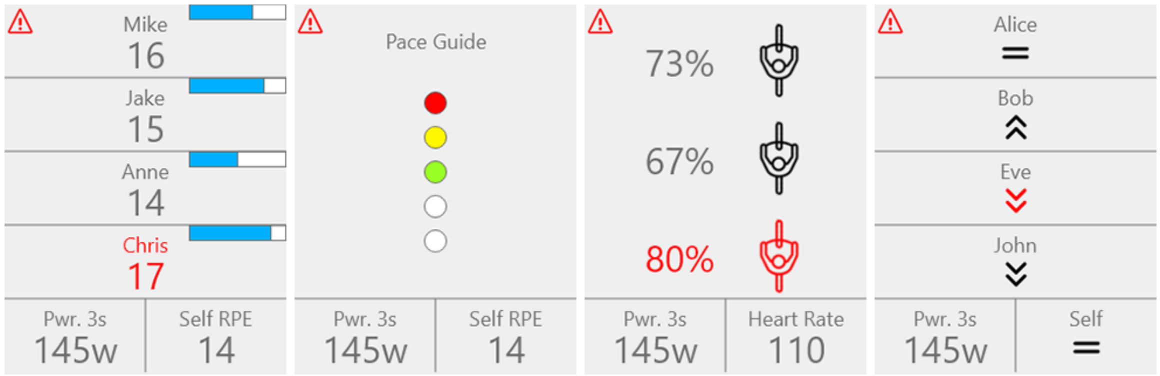

Figure 30 – Example of four bike computer (modelled after the Garmin Edge 530 screen space) pages sharing rider's effort. From the left to right: RPEView; Paceguide; %FTP, Requested Change in Pace. The red warning symbol on the top left is a quick indicator that at least one rider is above a pre-determined threshold.

I evaluated two designs in the study presented on this chapter: RPE View and Paceguide. RPE View provides a direct way for riders to learn about their peers' exertion levels but requires more time to read the data. Paceguide, on the other hand, gives an overview of the entire and requires little interaction time for the rider to read the information. The selected designs can be viewed from Shneiderman's visualization mantra (Shneiderman, 1996): Overview First (Paceguide) and Details on Demand (RPE View). Riders with access to both could use Paceguide as a quick at-glance tool and switch to RPE view when the need exertion information from a specific rider.

Both can use an underlying data source (e.g., %FTP, HR) to drive the exertion values shared with the other riders, however RPE was chosen as it translates across different riders with different fitness levels. In RPEView, riders can see their peers' individualized exertion values (in this case, RPE). In Paceguide, individualized values are hidden; instead, the layout computes a group-wise metric using individualized values and represents them as a sequence of colours (depicted in Figure 31) based on pre-determined thresholds.

Paceguide was designed with two concepts at its core: to have a small footprint, thus allowing it to be overlayed on top of other data fields (Figure 32), and to support at-glance visualization, minimizing the time riders must look away from the road. At the other end, RPEView focused on providing individualized information to identify riders that might be more tired than others, to keep them at the end of the paceline for a more extended period. It can also be used to identify riders that are more rested than others and therefore can lead (pull) the group for longer.

Figure 33 depicts a side-by-side progression of RPE View and Paceguide as seen by the participants in the study.

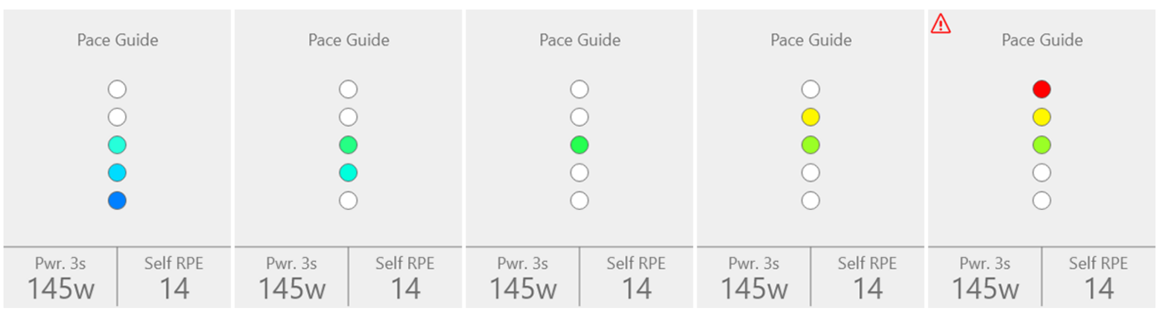

Figure 31 – The effort progression in the original Paceguide layout. Lit colours underneath the centre point represent low group effort (low training impact). In contrast, lit colours above the centre point represent high effort, with red meaning that a rider is likely to be dropped due to over-exertion. I opted to use a combination of "number of lights" and colour to facilitate at glance interaction as it allows riders to quickly look at the visualization and not have to count how many lights are lit (if they all had the same colour)

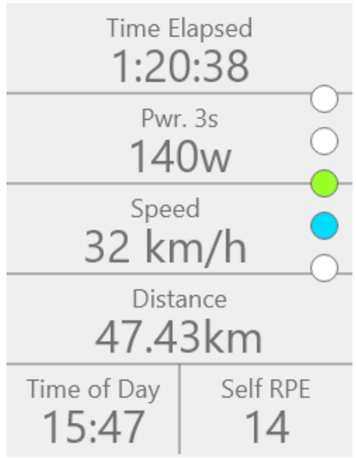

Figure 32 - Paceguide overlaid on top of a regular multi-field page

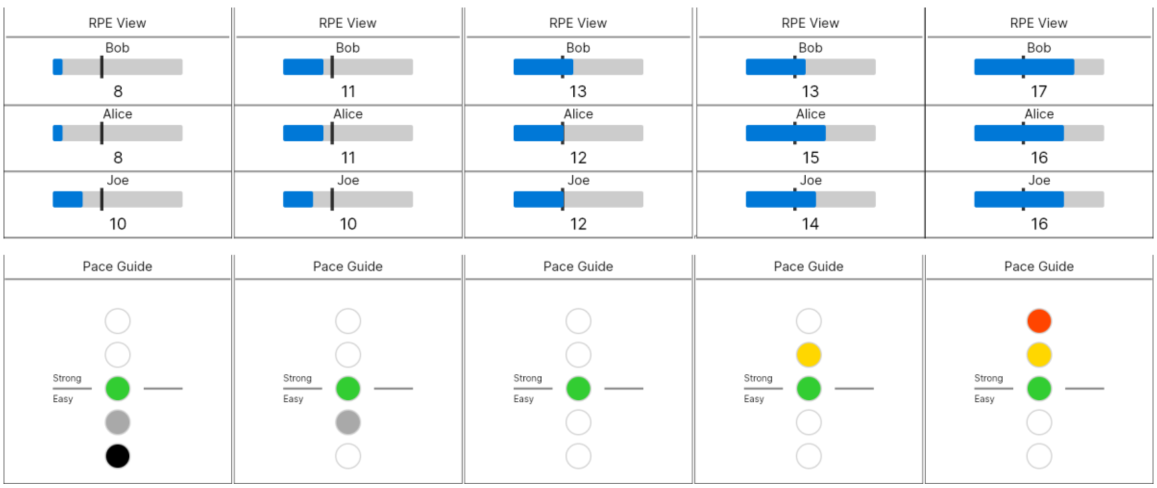

Figure 33 – Side by side example of the exertion progression on RPE View and Paceguide as seen by the participants in the study described in this chapter.

6.2 Considerations for the COVID-19 Pandemic

The initial design of the study consisted of participants riding physically together. Due to the COVID-19 protocols at the time, I adapted the study to allow group riding while keeping the riders physically distant from one another using the virtual group cycling platform Zwift.

I opted to use Zwift (*At Home Cycling & Running Virtual Training & Workout Game App*, n.d.), a virtual cycling platform, instead of virtual reality bike simulators like the one proposed by Hernández–Melgarejo (2020) so participants could participate in the study from their homes. Zwift is the current platform of choice of the International Cycling Union (UCI) for the esports world cycling tour. Zwift riders connect their bicycles to smart trainers, devices that can control the load on the rear wheel thus simulating changes in terrain (e.g., climbs) and air resistance due to drafting effects. The trainers can also sense the rider's power output, which drives the virtual avatar forward. Figure 34 shows the rider's view during the study.

Unfortunately, it was technologically challenging to use live data (e.g., current power output or heart rate) from the participants' smart trainers and sensors to drive the effort-sharing tools as Zwift does not have an open API that I could access. To use live data, I would have to develop my own group cycling simulator which was outside the scope of my thesis and not necessary for this research. Instead, I trained the participants on the RPE scale. I used the self-reported values instead, which meant that the researcher had to probe the participants' exertion levels at uniform intervals. There are advantages of using Borg's RPE scale. Because it is a self-reported measure, its values can be directly compared between riders with different fitness levels. Furthermore, RPE values can depict accumulated exertion over time, which is impossible with Watts, Watts/kg, or %FTP.

Figure 34 - The user interface of Zwift. Rider data is split between personal (top left) and ride (top centre). The study interface is placed on top of Zwift, covering the right part of the window

6.3 The Implementation of Paceguide and RPEView for desktops

The study application was developed in C# (.net (Core) version 5, Windows and Mac). It consisted of two parts: a control application used by the researcher to input the RPE values of each rider and an always-on-top application used by the participants to visualize the group data. Source code can be found at: https://github.com/jmfranz/AtHomeCyclingStudy The repository, however, will remain private until a paper of the study is published.

During the trials, participants would periodically inform their current RPE value (it is a self-reported measure) verbally so the researcher could update the values. The participant-facing application rendered three different screens depending on the study condition: baseline, RPEView, or Paceguide (Figure 35). In all conditions, the RPE scale was visible for the participants to consult. In the Paceguide view, the RPE thresholds for the green light was set to 13 (first ventilatory threshold, VT1) and above 17 for the red light (above the second ventilatory threshold, VT2) as described in (Alberton et al., 2013). To summarise the ventilatory thresholds, VT1 is when a person's exercise intensity starts to rise changing its breathing pattern and it is high enough to start accumulating lactate in the blood; VT2 is a higher intensity where breathing becomes very rapid, cannot be sustained for long periods and causes a rapid accumulation of lactate.

The group effort was computed using the RPE values of everyone in the group. If any rider was above 17, the red light was set to on, regardless of the RPE values of the rest of the group. If no rider was above the 17, the RPE values were averaged. The average value was then subtracted from the green line value (13) to compute the distance from the neutral RPE. This distance was then used to compute which lights should be on and off as per Figure 36.

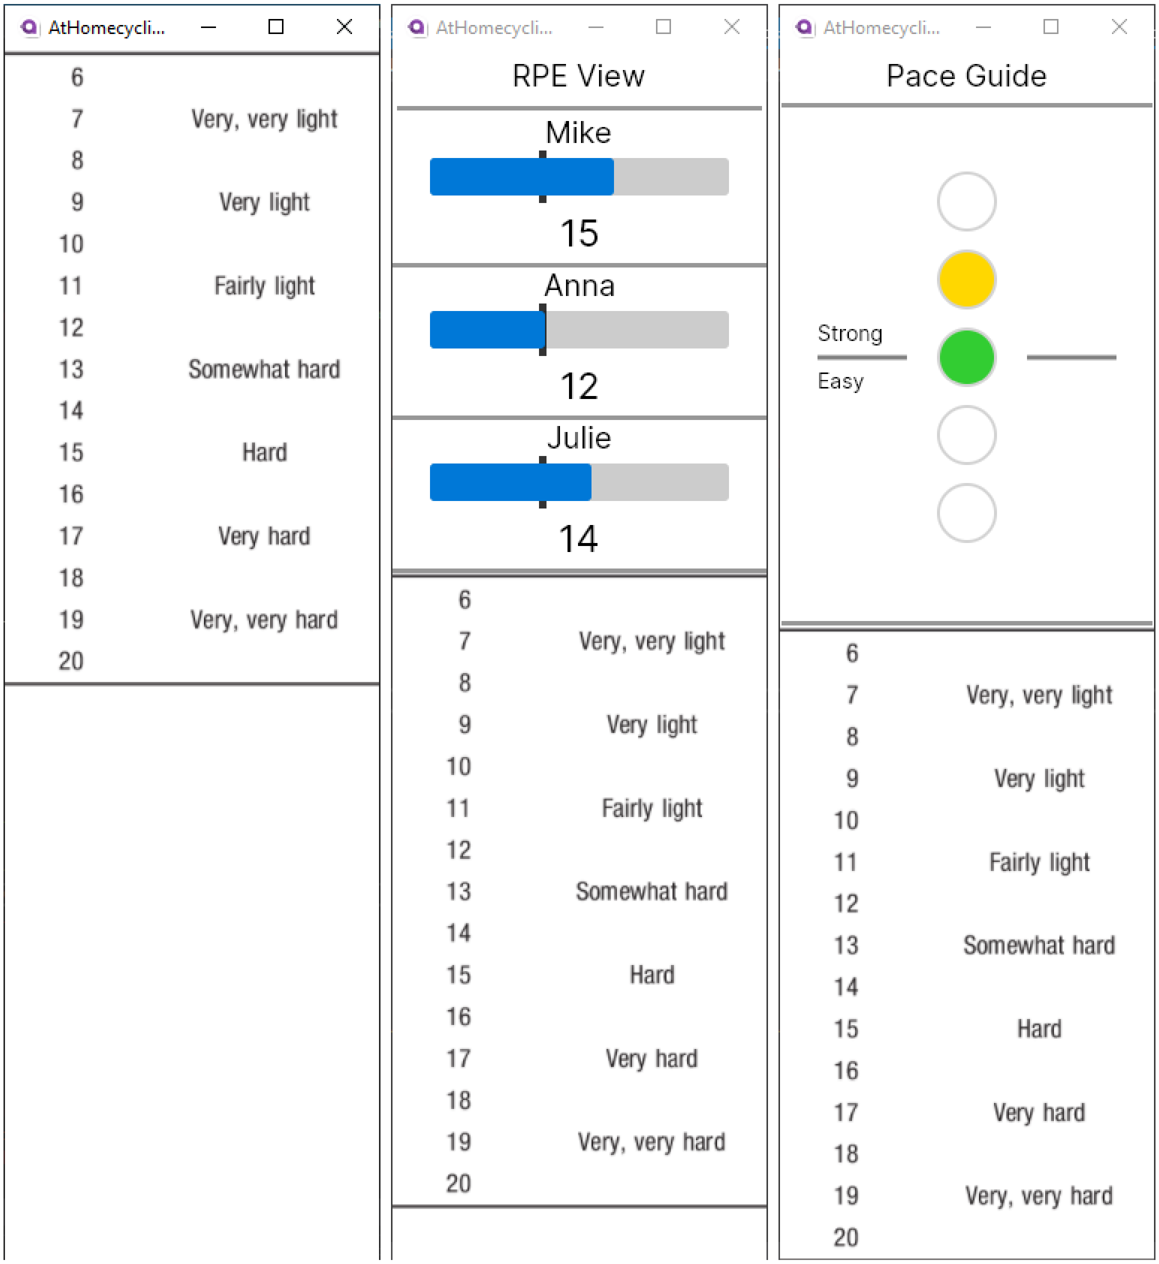

Figure 35 – The three views of the participant facing the study application. From the left to right, baseline, RPEView, Paceguide

From the computed group RPE, the lights were lit as shown in Figure 36. The middle greenlight was always on. The yellow light turns on at a positive distance from the NeutralRPE of one or more and the red light turns on at a positive distance from the

NeutralRPE of 3 or more. The values for the grey and black light are analogous in the negative direction.

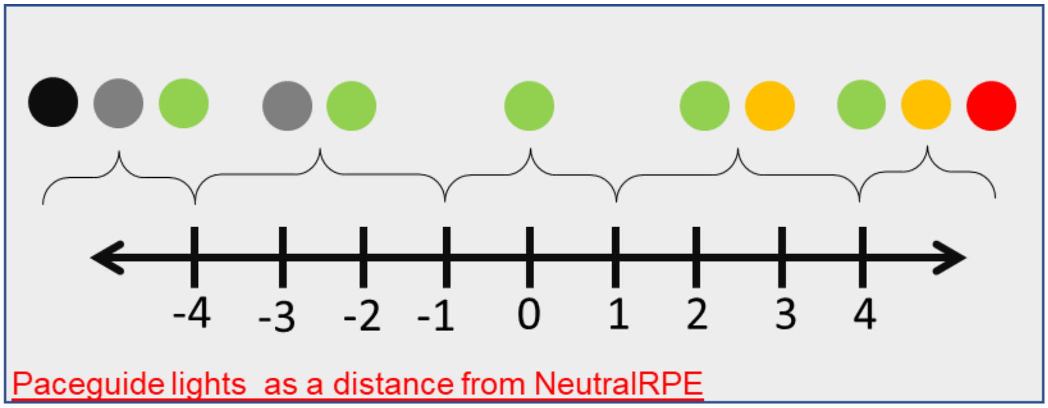

Figure 36 – Lights displayed on Paceguide as a distance from the NeutralRPE threshold (set as 13 for the study)

6.4 Methodology

To explore the research questions listed in Chapter 6, I designed a within-subjects repeated measures study with one factor and three levels: Baseline, and the two sharing techniques RPEView and Paceguide. In the study cyclists could ride together as a group while being exposed to the shared effort information while I measured the impacts on their performance in terms of power and number of drops from the group in a pre-determined course. I recruited participants in groups of three riders that rode together in the past.

The study consisted of five encounters (Figure 38) where three were group rides: one a baseline condition with no sharing, i.e., riding as they usually do, and two with sharing enabled. Following other studies in exercise physiology, participants always rode the baseline condition first and then I counterbalanced the order of the two sharing conditions. Furthermore, participants must have at least one day of rest between each trial and finish the study within two weeks to avoid performance differences between sessions due to their regular training sessions (Mieras et al., 2014).

6.4.1 Participation Criteria and Recruitment

To participate in the study, participants were required to ride at least 100km weekly and have experience with group riding in the real world and on the virtual cycling platform Zwift. Participants were allowed to use their weekly kilometres from previous years in case it was affected by the Coronavirus pandemic. Participants were required to be fluent in English, at least 16 years old and were not restricted to a specific region in the world.

Participants were recruited through social media, Today@Dal, and through a local cycling shop billboard in Halifax. Participants from the study presented in the previous chapter were not allowed to participate. I did not know any of the participants prior to this study.

6.4.2 Participation Protocol

Participants met with the researcher five times. In the first encounter, participants met the researcher individually to go over the consent form, give verbal consent to participate in the study, and become familiarized with the sharing tools and the study protocol. I opted for verbal consent for convenience; participation was remote and not everyone can sign PDFs. This reasoning was approved by the ethics committee.

After I ran the familiarisation protocol with the RPE scale with the participants during the first encounter:

1. Warmup up for as long as necessary
2. Ride "Fairly Light" for 2 minutes

3. Increase to "Somewhat Hard" for 2 minutes
4. Increase to "Hard" for 2 minutes
5. Increase to "Very Hard" for 30 seconds
6. Cool-down for 5 minutes

During each stage, I asked the participant if their feelings matched the cues for each keyword. The RPE table containing the association between keywords, values and body sensations can be found in Appendix M.

In the following three encounters, participants rode a pre-determined course of 20.1 km with the following instructions (a detailed step-wise protocol description for the rides can be found in Appendix P):

1. Ride at a strong training pace as a group
2. Rotate the lead rider every 5 minutes
3. Keep the group together (no drop ride)
4. Update the effort information every minute

I chose the Ocean Cliffside Loop (Figure 37) as it contains sections of flats, rolling hills and a climb while being compatible with the allocated time of each study session. The road conditions were mostly asphalt with a small section of gravel which is important as Zwift emulates the rolling resistance of the road surface.

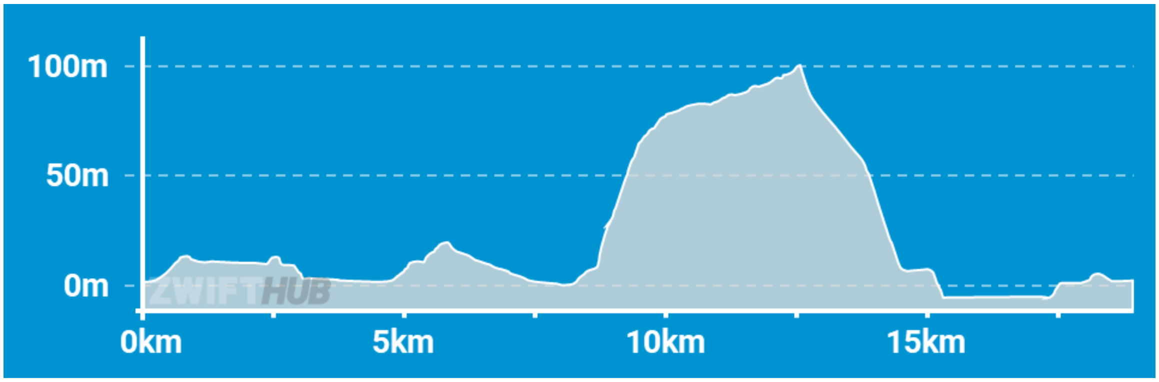

Figure 37 - Profile view of the study's course. Source: https://zwiftinsider.com/route/ocean-lava-cliffside-loop/

During each ride, I collected the ride metrics (power, speed, altitude, grade) for each participant through Zwift. I recorded the ride from a third-person perspective using Zwift's follow rider mode (e.g., Figure 34) while recording the screen on my computer and stored the progression of each rider's RPE values using the study control application.

The protocol for each ride was as follows: participants joined the researcher at a Zwift Meetup[22] five to ten minutes before the start time. Then the researcher reviewed the ride instructions and the sharing interface of the day with the participants. Participants would then start riding the meetup course forming a pace line. At 5-minute intervals, the researcher instructed the participants to rotate the lead rider. At 1-minute intervals, the researcher prompted the RPE values of each participant individually while temporarily blocking the other participants from hearing the spoken value. The information was then manually entered on the researcher's control interface and used to update the effort sharing view on the participant's computer accordingly.

[22] https://support.zwift.com/en_us/meetups-HJP7iUd4r

After each ride, participants answered a post-session survey with the following questions (one to five scale from Strongly Disagree and Strongly Agree):

A1. When I was leading, I was aware of my partners' effort level

A2. When I was leading, I was not sure how strong I could pull without dropping anyone

A3. When I was second, I knew how to strong to pull in my turn.

A4. When I was last, I felt that I was being overly taxed

A5. When I last, I was able to recover from my previous pull

A6. I have perceived the 20km ride during this session as

Assertions one through five had "1 – Strongly Disagree" and "5 – Strongly Agree" as the end points of the Likert scale while assertion six had "1 – Easy" and "5 – Hard" as the end points.

Participants completed the SUS survey after rides with the Paceguide and RPEView views.

After the three rides, participant groups met the researcher once more for a focus group discussion about their experience in the study and to reflect on the researcher's annotations from each ride. The discussion period focused on three main topics:

- Comparison between riding on Zwift vs. Real World
- Differences and impacts between the baseline and both tools
- Differences between the tools in terms of impacts, benefits, and drawbacks

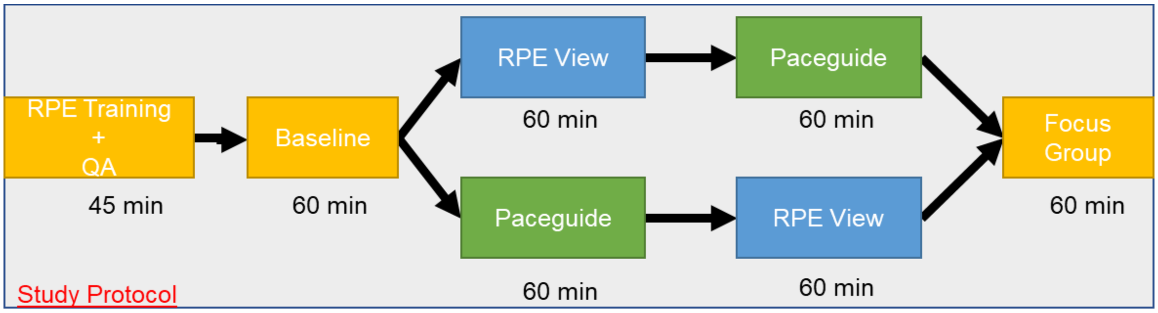

Figure 38 - Summary of the group riding protocol with the allocated time for each session

6.4.3 Instruments

The post-session survey was created from my interpretation of the work models created in Chapter 5. Through my contextual inquiry and work modeling, I identified key issues such as the lack of awareness of the group's exertion which my designs are trying to address. The questions presented in the post-session survey allow me to explore, as a self-reported measure, weather my designs were able to support the riders.

There was limited opportunity to fully validate the questions considering that they were used for a single study and not created with the goal to be used by others. The language used in the questions (such as *pull*, *recovery* and *dropping*) are terms commonly used in the sport and provide little room to erroneous interpretation.

Participants sport metrics' data for this study were measured directly through Zwift; the numbers of riders dropped from the group was coded from the video recordings, and the RPE value was manually inputted by the researcher during the trials. Table 11 describes the source for each variable.

Table 11 - Source of the data for the independent variables in the MMA study

Data Point	Source
Wall Time	Zwift
RPE Value	Manual Input From Voice Prompts
Performance Metrics (e.g., Power)	Zwift
Drop count	Video Coding

6.4.4 Data Analysis Tools

I processed ride metrics collected directly by Zwift (Flexible and Interoperable Data Transfer format[23]) in GoldenCheetah (*GoldenCheetah*, n.d.), an open-source cycling metrics analysis software. For each ride, I removed the first and last two kilometres to account for warm-up and organization times at the start of the ride and end-of-ride loss of focus, chit-chats and last-mile sprints. I then exported the individual processed data and performed the study analysis on R version 4.1. and the FSA package version 0.9.2. Audio data was first automatically transcribed using the meeting software and then manually reviewed to fix transcriptions errors. Video data was coded in NVivo 12 to identify the number of drops and the time it took for riders to recover from the drops.

6.5 Results

I recruited ten groups and collected a total of 8 groups of three riders. Two groups had to withdraw from the study, one due to COVID-19 and another from a race injury unrelated to the study. Participant ages were 44.52 ± 7.76 years; 14 participants identified as male and 11 as female. The mean Functional Threshold Power for the participants was

[23] https://developer.garmin.com/fit/overview/

159.42 ± 65.17 watts, while the average years of experience with group riding was 14 ± 12.47. Seven groups were from Canada (provinces of Quebec, New Brunswick, Ontario, and Alberta) and one group was from the UK. 22 participants used direct-drive smart trainers and two participants had fluid trainers. The power analysis presented in 6.5.1 did not consider the data of the two fluid trainers. Participants names used on quotes on this study are pseudonyms and to not reflect their real names.

6.5.1 Performance Metrics (RQ1 and RQ3)

I present the power comparison between the rides in terms of Normalized Power (NP) against the baseline rides (Figure 39). Allen et al. (2019) define NPs as *"a measure of the true physiological demands [or cost] of a training session*". Data was assessed for normality using Shapiro-Wilk ($p > 0.05$) and Bartlett's test for homoscedasticity ($p > 0.05$).

One-way ANOVA showed significant difference in NP (Figure 39) between the rides $F(2,45) = 10.759, p < 0.001$ and large effect size ($\eta^2 = 0.323$); post-hoc analysis with Bonferroni correction showed a significant difference between the baseline condition and RPEView ($p_{adjusted} < 0.001$). and between RPEView and Paceguide ($p_{adjusted} = 0.0057$).

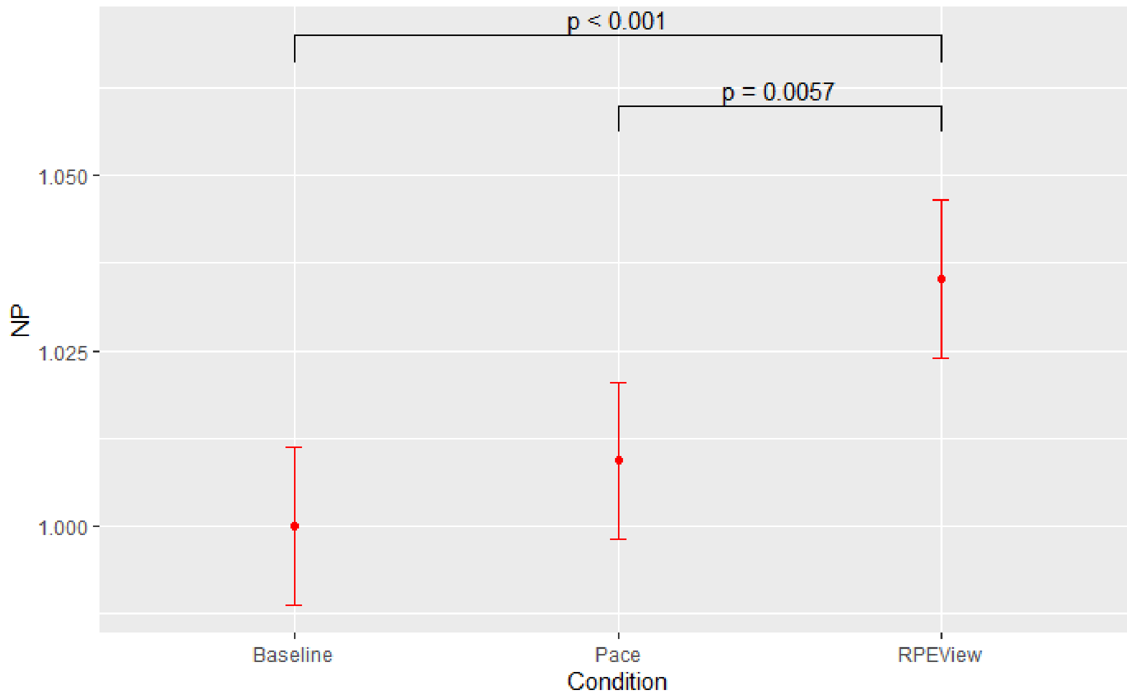

Figure 39 – Normalized Power as a ratio of the baseline value.

One-way ANOVA showed a significant reduction in the number of riders being dropped from the paceline (Figure 40) $F(2,15) = 12.917, p < 0.01,\ \eta^2 = 0.633$. Post-hoc analysis with Bonferroni correction showed a significant difference between the baseline condition and RPEView ($p_{adjusted} < 0.01$) and between the baseline condition and Paceguide ($p_{adjusted} = 0.0038$).

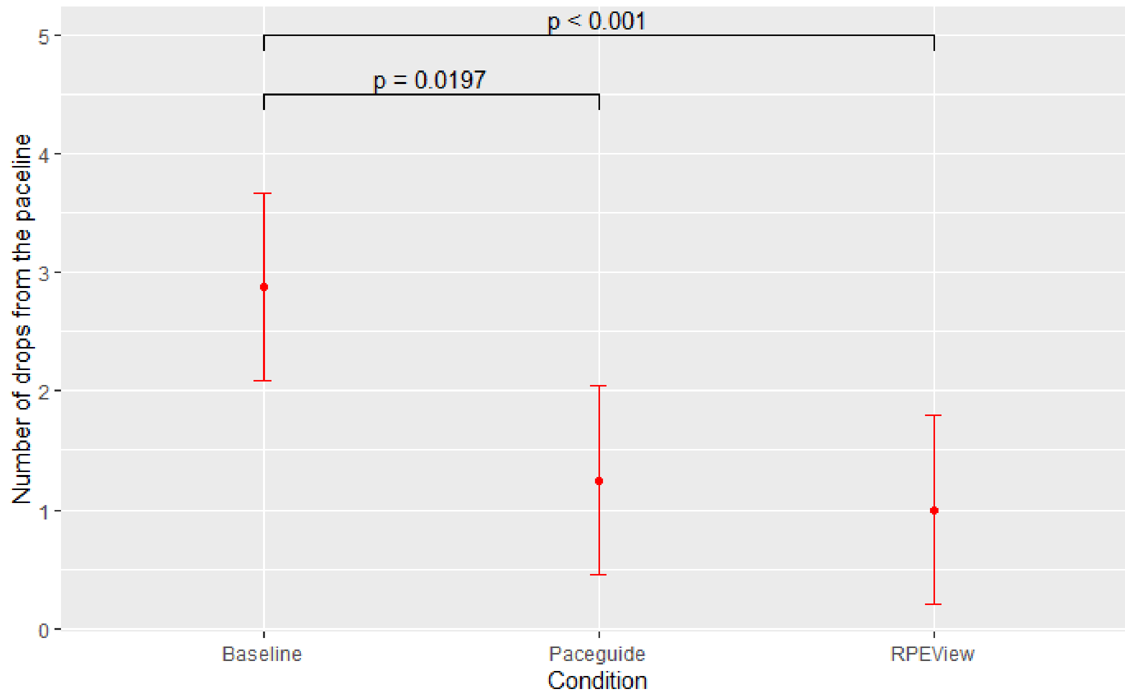

Figure 40 – Comparison of the number of drops between all three conditions. There was no difference between Paceguide and RPEView

6.5.2 Self-Reported Metrics (RQ2)

One-way ANOVA did not yield a significant difference for the System Usability Scale scores between RPEView and Paceguide; mean scores were 85.00 and 78.82, respectively.

One-way ANOVA showed a statistical difference in two of the post-session questions "*When I was leading, I was aware of my partners' effort level*" (Figure 41) $F(2,48) = 46.881, p < 0.001$, $\eta^2 = 0.661$ and "*When I was leading, I was **not** sure how strong I could pull without dropping anyone*" (Figure 42) $F(2,48) = 5.0818, p = 0.009$, $\eta^2 = 0.175$. For the first question, there was a difference between all factors ($p_{adjusted} < 0.01$) using Bonferroni correction. Participant data from both questions were not normally distributed; per Norman (2010) ANOVA tests are robust even when data is not normally distributed.

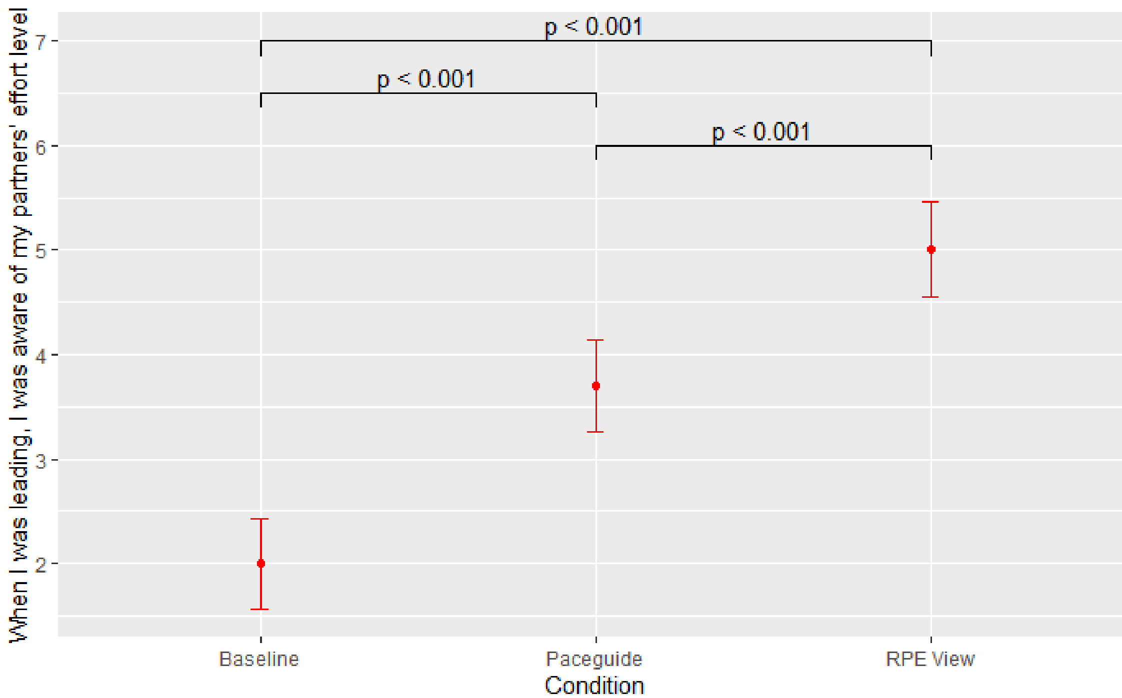

Figure 41 – Answers to the post-session question: "When I was leading, I was aware of my partners' effort level" using a 5-point Likert scale.

In the second question, there was a significant difference, in post-hoc analysis with Bonferroni correction, between RPEView and baseline ($p_{adjusted} < 0.01$).

One-way ANOVA did not yield a statistical difference in the other assertions: "When I was second, I knew how to strong to pull in my turn," "When I was last, I felt that I was being overly taxed," and "When I was last, I felt that I was being overly taxed," "I have perceived the 20km ride during this session as".

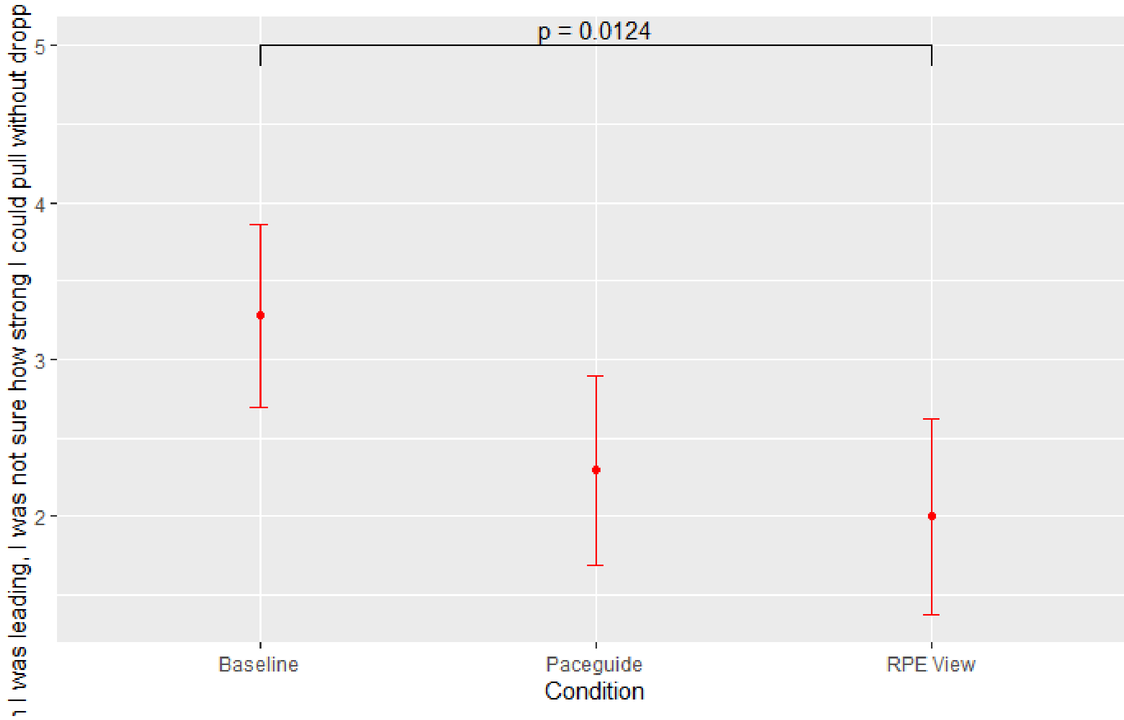

*Figure 42 - Answers to the post-session question: "When I was leading, I was **not** sure how strong I could pull without dropping anyone" using a 5-point Likert scale.*

6.5.3 Perceived differences between the baseline, Paceguide and RPE View

In this subsection, I present the summary of the focus group discussions with the participants after they were exposed to all study conditions. I used the issues identified in the work modeling from Chapter 5, such as lack of awareness and difficulty pacing, as a theorical foundation to aggregate the quotes and comments from the discussion period. Audio data was transcribed to text and then comments were tagged using the aforementioned pre-definitions using Nvivo 12.

Participants unanimously agreed that the tools helped them pace the group better and keep everyone together. They mentioned using the tools to validate their pacing decisions when leading. The tools made their task more straightforward than just relying on their perception of the others.

When I am leading, I will never just blow-up, go beyond Zone 3 [Power Zone], but the numbers helped to validate: okay, no one's blown up or no one disappeared or is suffering greatly. – Alice

When you're on the front and you could see people putting out similar efforts to you, I think it definitely helped you because you know the group is all working in a similar level and quite often you know one of us on the front might be a 16 or something and the others were 14 or 15. Then the overall level is; this is a good pace to maintain. – Bob

When I saw the number going up for someone, I knew I had to slow down. Without it, I was constantly checking if someone was dropping, I had to figure it out by myself. I think the two views helped a lost, we had a better cohesion. – Eve

Riders agreed that they had better awareness of their peers and that the shared effort information would be even more beneficial in larger groups where the leader has even less awareness of the entire pack.

If you are on a long paceline and you are pulling, you can't see what the last person is doing. So, if our goal is to stay together as much as possible it helps to know how much effort everyone is because we do drop people through our rides and don't realize until we stop at a traffic light and realize half the group is gone. – Charlie

I've found that with the Paceguide and the RPE I was more aware of Karl and Norton – Frank.

Three groups mentioned that access to the effort information helped them see how hard one can push the group without anyone dropping. The concept of trying to push the envelope might have come from the ride instructions *"Ride at a strong training pace; you* *want to get some training impact out of our session."* Participant Dave summarizes the concept:

> *You want to make people suffer but just enough, so they don't drop. I am a bit of narcissist, so I liked to know how people are doing. If they are not working hard enough, I will keep pushing the pace. I loved to have access to the effort of others. I guess it translates to what you do in the real world too as you are always pursing that power or heart rate value. – Garret*

When comparing RPEView vs Paceguide, participants leaned more towards RPEView (3 participants preferred Paceguide). They described it as more accurate and less "*open to interpretation*" [Alice], with less "*guessing work*"[Karl] of who is having a hard time keeping up. With RPEView, there was no need to "*recall the effort range of each coloured light*" [Eve]. Participants from three groups mentioned that they were able to combine the individualized values and their internal knowledge of their peers to have a clear picture of the group.

> *I like the individual numbers; in a group the traffic light is difficult then to determine kind of minor adjustments that needs to be made for a given rider that is suffering – Harry*

> *The numbers were better for me. I could individually view each person's number and knowing Uri's and Lam's strengths I can kind of put my mind inside their mind and know exactly what I need to do for us to stay together as much as possible. – Jaimie*

> *I prefer the RPE View because after riding with people for a long period of time you kind of know who might get dropped. Therefore, I could use it to keep track of those riders as we go by. – Karl*

With the RPE view you have cleared information. You can see who is struggling and can plan on how to help that person while you are doing the rotations. – Lam

Furthermore, four participants mentioned that it was hard, if not impossible, to understand who was driving the changes in Paceguide as that information was unavailable. This could lead to doubts on how to manage the situation to increase, decrease or keep the same pace.

I was more conscious with the RPE. With the traffic light [Paceguide], I was not sure who was driving the lights. If I saw a yellow, I wondered if it was me or someone else. With the individualized numbers I could see that I was the one driving higher values for whatever reason. – Mike

I was staring at the lights and I just wondering, based on what I was feeling and the lights, what are the others feeling. The only thing I could gauge is that independently of who is causing the changes in the lights, if one of these strong riders is struggling, I must slow down, doesn't matter who it is. However, with the RPE it was easier for me to stay in the paceline better. From the numbers, I knew if Lam had no more matches or if Uri had the potential to speed up, I knew approximately where I was and what I had to do to maintain my position based on their energy levels. I still find that having the numbers associated with their names helped to make some educated assumptions based on my previous experiences riding with them. – Jaimie

Although there was a general preference for RPEView, three participants argued that Paceguide was better because it provided a way to read the information quickly. The coloured lights made it possible to glance at the tool and understand the overall group metric instead of thinking about the individual RPE values.

When you are working hard on the front, you just want a quick snapshot of the data so the traffic lights kind of work towards that sense. – Olaf

Less thinking is better so probably the lights – Norton

I've found the Paceguide more intuitive. I was interested in the numbers, but it took more thought to figure out what the information was there. Maybe if the order changed [name orders] to reflect the position it the paceline it might have helped – Paul

Two participants pointed out that the choice between group metrics versus individualized values must consider how homogenous the rider's fitness levels are. In a heterogenous group, the group value might not represent the actual exertion state due to individual differences. I have observed this phenomenon in the study where one of the groups was going up the climb section with only the green light one, but one of the riders was putting a higher effort than the others. Closer to the end of the climb, another rider raised their RPE value by one, and the light switched to green-yellow: "*Finally someone is suffering with me!*" [Tessa].

Both Karl and Norton have valid points. I assume if everyone is around the same fitness level, traffic lights might be better. But if the group is not well balanced then the individual numbers are more helpful. – Frank

I don't think the traffic lights paints a very clear picture. When I see those hills coming and I am just dying, and the light is yellow. I know the other's strengths so am I the yellow and they are green or am I close to read and they are yellow? And I think that as the group grows, then these individuals that might be more tired for whatever reason are going to get lost in the data. – Jaimie

In the focus group, I asked about possible privacy concerns of showing individual values to everyone and opening the possibility of exposing riders in a group that might be underperforming or "unfit" for the group. All participants believe that that was not an issue in groups that regularly ride together, and the benefits of being able to identify weaker/tired riders outweigh any privacy concerns. Participants mentioned that, in some form, riders that exercise together have an idea of the fitness levels of others either from the experience of riding together or from sports metrics shared on Strava or other sport-focused social media.

> *In a group setting, if someone is suffering and you can identify them early, you can take care of that person making sure that he or she is sitting more in the back while less tired riders can spend more time in the front – Mike*

> *I think if there is a rider that is tired because they overtrained or had a heavy night, knowing it is useful. Because then there are not what you would expect them, you know. So, you get extra insight into how they're how they're really doing. – Paul*

> *Yeah, we've all had situations where you are just overtraining or feeling ill or coming off a cold and you go out for a ride, and you are just struggling thinking how this is going to be tough. I think communicating that is a good thing because otherwise people may, particularly if you do know the dynamics of that group, be expecting you to do those extra 100 watts that they normally do. If they haven't got them and then I think you need to let people know quite quickly, otherwise they might think you're just freewheeling or being a bit of a wheel sucker. – Garret*

> *In groups that know each other, you already know who's going to drop when there is a climb. It would be nice to know, just before you start*

> *climbing, what is this person state. Because if you are too fast, even on the flats they are probably already taxed. I liked knowing the state of the others so we can adjust. - Quinn*

When questioned about their thoughts on sharing exertion information in a group of strangers, e.g., an ad-hoc ride, all participants agreed that it would be beneficial to fill the knowledge gap of the fitness of others. They mentioned that it could help them pace the group when they are leading without pulling too strong but also not pulling too easy.

> *Assuming that everyone is using it and I want to keep a pace where the group feels happy. I want to know that everyone is together and is at a OK pace so I'm not going too hard and drop someone. That's my feeling so I want to look at that happy whatever that color is whether it's green or whatever and that's. – Frank*

> *I think it would also help me to understand how long I should pull for. If it's if the group is way stronger than me then I would pull hard and take a really short pull, right? But if I knew I was one of the stronger riders and I felt fine but everyone else was getting tired, then I would pull for longer and go slower or go as hard. – Karl*

> *It would help you. Because we're coming at this from a point of view of having that background knowledge that we've just talked about and having an awareness of basic fitness levels, basic ability, basic power. – Garret*

6.5.4 Ecological Validity and Limitations of Zwift

The sharing effort tools I proposed in this study target riders in the wild. However, due to the COVID-19 pandemic, conducting the study with riders physically together was not feasible. As a compromise, I opted to use Zwift as the study's cycling platform.

Therefore, it is essential to address the limitations of the software, considering that it is a simulation. During the focus group discussion period, I addressed how participants perceived the differences and similarities between riding on Zwift and outside. From a research perspective, however, using a virtual cycling platform allows for better control of confounding variables such as weather conditions, road surfaces, and car traffic. At the same time, differences in rest states and the performance between different bicycles are still reflected in the simulation.

All but one participant believes that their experience riding inside reflects their expectations of outdoor rides except for some key differences. Participants indicated that the power output to keep up with the virtual group was close to their output on rides during the outdoor season.

> *My effort is always judged on how much effort I'm putting on versus what I'm seeing behind me or in front of me depending on where I am. I don't know whether it's any different than from riding outside. - Riley*

> *I do [believe my experience on Zwift translate to real world rides], reasonably. All I would say having ridden with these two guys doing in the real world is that it was a fairly similar experience. I was hanging on for dear life most of the time, so it was quite close actually. – Paul*

> *I would think that the [power/exertion] feel is very it's similar for me, I think. You know, I think it translates well. I wouldn't say it's exact. I think riding outdoors is still easier. – Jaimie*

> *I think it does [my experience on Zwift translate to real world rides]. I ride with Riley quite a bit and so the comparison between inside versus in real life is close. – Silas*

Regarding power, I've done FTP tests on Zwift and out on the road and I think they're pretty consistent. – Garret

I think it's as Paul would say reasonably close but it's certainly not an exact match. – Olaf

The effort inside and outside is pretty much the same. I think outside can be more difficult depending on the road conditions [e.g., new vs old asphalt] – Eve

I think on Zwift with reality is the group effect with drafting is simulated and it's somewhat close. If you get dropped, you feel that it's hard to get back on, and then when you're on and you're in it, it's pretty good. – Frank

I think it can check a lot of boxes, but it's never going to completely replicate what you do outside in the road because you haven't got all those variables of traffic and road surface and wind direction and those kinds of things. – Bob

Participants also mentioned that their perception of the effort they had to put in during the study and the RPE values they have seen of the other riders in the RPE View condition matched their expectations from their internal model of their peers' fitness levels.

I think it match up well. I think it match mostly for me, especially on the last day when we could see the RPE values. – Mike

I agree with Mike. I know Mike and Silas are both stronger than me and so I like knew that, probably most of the time it would be to keep up with them and I think that was true, especially the first two rides. – Riley

However, participants mentioned that riding inside was relatively harder than outside for the same amount of time due to the lack of coasting, stops (such as intersections and traffic lights) and the bike's rigidity on the fixed trainer.

> *You freewheel [coast] more [outside] so I feel more tired after one hour Zwifting than I do after an hour out in the real world, but I probably have put in more kilo joules overall. – Rob*
>
> *I think 2 hours on Zwift is worth 3 hours out in the road because of the amount of freewheeling, stopping for traffic, stopping injunctions, and things like that. – Garret*
>
> *One of the biggest things when we ride together outside, we hit traffic lights; we start, we take breaks for chatting. So as the main difference, I think when I'm riding one Zwift, it's always much harder effort then outside. Again, for our group there's a lot of breaks and that personally works well for me. – Lam*
>
> *It is shorter and more intense. You could go out for a three-hour ride outside and it is okay; but if you do a three-hour ride on Zwift, you will feel it. Maybe it is just because you sat on your bike all the time and it is not as interesting. It is easier to do long rides outside. It is just like running 20km on a treadmill or running 20km outside; it is a different sensation – Tessa*
>
> *The bike doesn't move, it feels stiff whereas outside it feels like liberty – Eve*
>
> *I think that the integration of your body on the bike and [lack of bike] movement. I can ride for 4 – 5 hours outdoors but inside anything more than 90 minutes makes me feel saddle pain. - Jaimie*

The same goes for engagement and motivation to push harder on longer rides.

> *I feel like there's a much more engaging experience no matter what when I ride outside, I feel like I've ridden a lot harder and gotten a much better work out. – Karl*

> *In some sense, I find it easier to push myself out on the road because you are not sitting in your shed or in your house or whatever, and it's easy to kind of motivate yourself - Garret*

6.6 Discussion

As shown by the previous study, riders faced difficulties trying to gauge the current exertion of their fellow cyclists while trying to pace the group without dropping anyone. The study described in this chapter presented the same issue in the baseline condition's self-reported effort awareness results and in the discussion period. Furthermore, a significant difference in the number of riders being dropped also reflects the difficulty of pacing the group while maintaining cohesion.

Results from the study show that both interventions could mitigate the lack of effort information with varying degrees of success. Paceguide and RPE View enabled the riders to understand their peers' exertion better. However, RPE View was more effective in supporting the leading task as observed by the very high confidence answers to the *"When I was leading, I was aware of my partners' effort level"* question and increased performance while not dropping riders.

As observed in the results, participants had a significant increase in Normalized Power while using RPE View. It shows that riders could ride stronger than they initially thought without hurting the group. While one might argue that a four percent change might not have a significant impact overall, as ride distance increases, the difference in power output will lead to shorter ride times. The study's route was planned with a constraint to keep each session within a one-hour window; however, typical ride distances are longer than the 20km loop selected. When riding longer or for longer periods of time, effort management is of the utmost importance to ensure riders can finish the course. Therefore, riders who know their peers' effort and can use it to pace the group have an advantage over others.

While I did not observe the same NP difference in Paceguide, it should not be discarded: it can be considered safer because it requires less time to gather information and it uses less screen space. Furthermore, Paceguide scales to larger groups better as it requires a constant screen space regardless of how many riders are in the group. RPEView would require scrolling due to the limited screen space in the bike computers. However, further work is required to refine the algorithm that computes the group effort with larger groups, so an individual rider is not underrepresented in a large pool of riders, e.g., by using a weighted average instead of the current distance measure.

Similar to Walmink's work (2014), one of my goals was to increase the mutual awareness of the rider's exertion levels in the group thus supporting better informed pace and paceline positioning decisions. Self-reported exertion values (RPE) are an excellent proxy for measuring and sharing exertion as they translate to different riders independently of their fitness level (a 16 RPE value feels the same for different people) in contrast to absolute HR values used in (Walmink, Wilde, et al., 2014). However, the trade-off of RPE is that riders must manually update their current exertion state often. This manual data entry works in a research environment where the researcher conducting the study is constantly asking for updates. However, it might not transfer to riders riding in the wild.

Other exertion metrics should be considered in the future. Relative metrics such as current power zone and %FTP are a good start since cyclists already have power information while riding. One issue, however, is that power metrics do not consider the time riders have been riding at a specific zone. For example, as an athlete keeps exercising at constant power, it is expected for stress to accumulate even though the power zone and %FTP remains the same. When using RPE, one would expect the value to go up over time, depicting the increase in stress over time even at a constant power output. %HR and HR zones, on the other hand, can provide meaningful information for long rides as HR will increase over time; however, there are other issues with the metric as previously discussed.

An automated or semi-automated method for computing exertion overtime is needed to support group riding while not adding another high-demand task to the riders. One approach that is work considering is to explore Skiba's W' concept to measure how tired the riders are throughout the ride. W' closely relates to the idea of a battery, and riders can deplete and recharge as they change exercise intensity. The concept of a battery is already reflected in the W'bal (W' balance (Skiba & Clarke, 2021)) used in the post-ride analysis software Golden Cheetah. While the metric is not flawless and requires a specific testing protocol to estimate the value, it is the closest to an automatic metric that relies on the sensed data in the cyclists' computers.

Another possible path is to explore new metrics in conjunction with those created by sports researchers. As shown, RPE works well and can reflect influences outside the numbers such as tiredness from previous training sessions or bad sleep nights, but it is a manual process. Future effort-sharing iterations could leverage computer intelligence to use the sensor information to model the rider exertion over time and prompt the rider to update their RPE value when it detects changes in pace or combinations of value not reflected in the model. Machine learning could, in theory, model the entire energy expenditure system and try to infer effort values based on the sensor data. An exciting research area is how to allow athletes to modify the prediction model to fine-tune it for specific rides, e.g., tired and want to go easier.

6.6.1 Paceguide and groups with different fitness levels

One problem I identified with the Paceguide was its inability to reflect changes when only one rider pushes stronger than the others. This situation occurs when a rider is weaker than others (or is tired from previous training sessions). When a tired or weaker rider is not represented in the group effort, the group cannot support this rider leading to more exhaustion.

6.6.2 Privacy is secondary

One concern of sharing personal information is the impact it might have on the individual sharing the data. In the case of the riders, there was a question if, by sharing exertion values, we could be identifying weaker riders that could be slowing down the group and thus generating attrition between the riders in a group.

Results from the discussion period refuted the initial concern relating to Nissenbaum's work (2004) suggesting that the concept of private information and its impacts on sharing is anchored in the current context of the people sharing.

In the group cycling context, we can see the advantages of sharing personal effort information and using it to identify riders that are weaker in the group or just underperforming on a given day. Among riders that know each other, knowing who is struggling can aid weaker riders by allowing them to skip the leading role when rotating, thus protecting them. In groups of riders that know each other, knowing the other riders' exertion levels and fitness levels allows whoever is leading to have more informed pace decisions promoting a more uniform ride with fewer drops.

6.6.3 Going Virtual

One possible limitation of my study is that riders were not physically co-located and were riding at their home using Zwift as a virtual proxy to real-world group riding. A strong psychological influence when athletes are exercising together (either as a team or individually but co-located) is the Köhler effect or Köhler Motivation Gain Effect (Kerr & Hertel, 2011) (Original article in German (Köhler, 1927)). Köhler and the subsequent studies demonstrated that athletes with different fitness levels working together benefit from group work and perform better than if they were working alone, especially the weaker ones in the group. The effect exists even when two athletes that do not know each other are exercising together (co-active) – like two runners sharing the same route by chance. However, it is stronger in groups working together where their performance depends on the weakest athlete (conjunctive), like group road cycling[24]. According to the work by Kerr and Hertel, to better support the Köhler effect, some steps can be taken (among others):

- "Increase the likelihood of social comparison by providing individual and comparable performance feedback to the group members."
- "Making individuals and their contributions identifiable."

[24] The tradition example is if mountaineers climbing tethered to each other. The group can climb only as fast as the slowest climber.

- Increase the likelihood that inferior members perceive their efforts as crucial for the team's success, for instance, by implementing conjunctive task demands
- Making it easy for teammates to monitor and evaluate each others' performances continuously

Moving from physically collocated to remote cycling in the present study did not interfere with the points listed above. Riders were still working together, could hear each other (including heavy breathing noises) and see their partners' avatars[25] changing position on the road (e.g., moving away when dropping or moving slowly to the back of the paceline when rotating). I argue that there was an increased support for the Gains Effect in the study, despite being virtual, since the two study conditions introduced new tools to continuously monitor the rider's performance while riding.

From a sports physiology perspective (instead of sports psychology above), transitioning to indoor cycling did impact the riders' performance. As mentioned by the participants, riding indoors feels more challenging than outside. This sensation of "being harder" has been explored by Mieras et al. (2014). The authors demonstrated that when given free ride instructions – like those in my study: "ride at a strong training pace" – riders performed up to 30% better outside than in laboratory settings at the same RPE levels.

Therefore, considering Mieras's findings, the performance gain observed in the RPE View can be even more significant if the tools are used outdoors instead of in a laboratory setting. Furthermore, running the study on Zwift allowed for greater control of confound variables

[25] Irvwin et al. demonstrated that virtual avatars are enough to promote Köhler's Motivation Gain Effect (Irwin et al., 2012)

(such as weather and traffic) without nullifying the transition of the findings to outdoor riders.

CHAPTER 7 DISCUSSION

Through my work, I focused on designing and applying tools to promote awareness and coordination in collocated groups enabling users to see role specific information while maintaining some form of common reference available to everyone to facilitate engagement. I have considered different roles access to information from different perspectives. Sharing information while providing context to fill knowledge gaps between multiple parties plays a critical role in promoting engaging experiences independently of the collaborative nature of the experience.

From the study at the Maritime Museum of the Atlantic, we can observe that symmetrical, i.e., everyone sees the same thing and interacts with it the same way at the same time, is not necessarily required. Asymmetrical (and asynchronous at times when "look ahead" exploration emerged) sharing – people seeing different aspects of the data or interacting with it differently from each other – can promote interaction as long as there is enough to ground conversation and engagement between viewers; provided by the Halifax Harbour wood model and the Indicator Rings (when available). This finding is interesting as previous museum technologies focus on providing the same experience for visitors exploring the museum together; audio guides, hand-held AR devices, video walls and static signage all provide the same content at the same time however my work demonstrates that museums can explore new content avenues where visitors are exposed to experiences from different perspectives while still interacting with each other.

The success of the collaborative experience using the Semantic Linking as a "*What you see is not what I see*" (Snowdon et al., 1995) approach is on-par with Zhu's work (Zhu, 2004); however, as demonstrated an explicit definition of roles before the starting the WYSINWIS collaboration was not required in a content exploration/learning experience scenario such as the museum exhibit. Furthermore, with the introduction of the Indicator Rings, I have shown that it is possible to blur the boundaries between WISIWYS and WISINWYS; with it both participants had parts of their interaction connected (WISIWYS) through the rings and we even observed informal role switching e.g., the participant not wearing the HoloLens using the rings to "take control" of the experience while asking for specific buildings to be selected.

As observed in the museum, all sharing methodologies promoted some form of engagement – even the minimalistic Indicator Rings. Each technique had advantages and drawbacks (e.g., the higher mental load on Over-the-Shoulder); thus, choosing the ideal method for future applications should be considered on a case-by-case basis. Approaches similar to Over-the-Shoulder sharing work well in situations where pairs have similar interests (e.g., explore the 3D models fully, read the signage at the same pace) however it does not scale well to larger groups, such as museum tours, or when the pairs have different paces or goals exploring museum pieces (e.g., one focus on the visual aspect while the other's goals are to read and interpret the textual information).

Results from the MMA work informed the design choices made in the following studies. Like riders, each participant in the pair had a different role when interacting with the exhibit and its virtual content, albeit riders have fluid but well-defined roles[26]. The design focus of

---

[26] I expect that outside a controlled study environment, museum visitors would take turns in using the headset thus making role changing more flexible.

my work was to rely on existing tools already in place in the workspace of cyclists, i.e., the bicycle computer, and not "replace existing mediators and artefacts" (Bødker, 2006).

The interaction structure in the museum does not translate well for cyclists. Not only is the technology different (head-mounted displays vs. heads-down displays), but their goals are essentially different. Albeit both are working together, museum-goers are looking into exploring the space at a relaxed paced while exercise minded riders focus on finishing their training session while gaining some training impact from it. Furthermore the physical demand context in which they are interacting with digital systems are on opposite edges of the exertion spectrum.

Results from the MMA work inspired me to explore information sharing between people while they are exercising as a group with different roles supporting the work like the maritime information AR tool and the museum visitors. I have directed the last stage of my Ph.D. work towards road cycling. It begins with understating the work done by riders while they are exercising. While I had anecdotal evidence of some of the issues encountered by road cyclists, a deeper exploration of their work and needs was required to equip future research and focus the design on supporting them.

The Contextual Inquiries showed that riders have a well-established understanding (or frame (Orlikowski & Gash, 1994)) on how their exercise technology and information systems (i.e., the ubiquitous bicycle computer) support their work and how they can leverage them while training or during cycling. The sport cycling community's technological frame around the computer is mature enough to enable technological use of devices across different manufacturers. While technology has evolved since the original mechanical computer to care for the needs of an individual rider, there is no support for between rider communication and sharing information within the same well know technological frame. Participants were excited by the "idea of sharing" but were so set on their understanding on the capabilities of their devices that the concept of seeing exertion or biometrics data from the other riders was foreign to them.

The findings also showed that while researchers are exploring cyclists commuting or exploring the city as tourists, road cyclists have different needs, and most of the findings from the literature do not apply to them. Because there is a well-defined technological scope for road cyclists, I argue that research with alternative forms of feedback such as vibration (Poppinga et al., 2009; Vo et al., 2021) and projection surfaces (Dancu, Vechev, Ayça Ünlüer, et al., 2015; Kadomura et al., 2014) might be seen (from a sport perspective) as toys, adding weight or taking precious real space on the handlebars. A notable exception is Walmink's heart rate sharing helmet (Walmink, Wilde, et al., 2014) for pairs of cyclists. Although it does not scale to groups and uses absolute HR as the sharing metric, it demonstrates the feasibility and hints at the impacts (although from a social aspect) of sharing bio-signals among riders. My work improves from Walmink's with a proper evaluation from a sport performance perspective using meaningful bio-signals metrics, with sharing mechanisms that support multiple riders and focusing on the flow of effort information from back to front (most meaningful) as from back to from (as Walmink's).

Since group cyclists are riding with others, some roles and responsibilities go beyond the individual (as opposed to a commuter or child rider (Matviienko & Ali, 2019), for example). As we can see from the results, three roles exist, the leader, middle, and wheel.

Each role comes with its unique set of tasks and cognitive load as expected from a traditional work environment like a shop manager or a kitchen's cook. Interestingly, unlike most office work, roles are fluid, and riders switch roles often as they move throughout the paceline during their rotations.

This unique characteristic poses an interesting engineering technical challenge: to best support riders, one should display only the information relevant to their current role. It saves screen space in the already small displays and increases safety as riders will have to spend less time looking away from the road. However, identifying the rider's current role as they move through the paceline is a problem that needs to be addressed before role-based information filtering is possible.

Between riders, communication was another challenge exposed by the contextual inquiry. It directly impacted the rider's ability to collaborate with their peers. Communication is not only affected by environmental aspects like wind, rain, and cold, which causes "data packages" to get dropped from time to time, but also essential information such as the exertion state of riders is not shared among them. Information about their bodies and physical state is limited to instantaneous measures such as HR and Power. In contrast, information such as accumulated exertion, resting state, and burning sensations are not collected nor shared.

Riders, athletes of other team sports, and coaches can benefit from more work on how to best collect implicit exertion information and share it with their peers while actively engaging in sport (thus reducing the social-technical gap (Ackerman, 2000)). While it would not completely solve the pacing issue in cycling, reducing the knowledge gap between the rider’s mental model of their peers' state facilitates the optimization problem of pacing a group.

> *"So, in the end, you cannot really try to optimize the pace. You are just making sure you are within a threshold, that you are not killing anybody. Even if it costs not being at the optimal pace. I am just trying to keep the group together".*

Riders exposed other aspects of their work during the Contextual Inquiries, described in Chapter 5, such as problems with their bike computers, gestures, safety, and work environment. While the Contextual Inquiry study identified multiple avenues for future work, I focus on how to share effort information among riders and better support the leader’s pacing decisions. Correct pacing plays a significant role in the success of a group ride as bad decisions can lead to group fragmentation, frustrated riders, and slower times.

The final study (Chapter 6) explored the impacts of two different exertion-sharing tools (RPE View and Paceguide) on group cyclists. Both tools were designed to be compatible with bike computers as they are ubiquitous among enthusiasts and professional riders and very popular with recreational ones.

While there was exertion information available to the current leader, that data was only related to their own body. Leaders had to use their data to try to guess or estimate the group's overall effort and make pace decisions based on missing information. That leads to situations where the group is riding at lower intensities or, worse, higher ones, thus leading to riders being overly taxed and dropping from the group.

Results from the study show that both interventions mitigated the lack of effort information with various degrees of success. By introducing ways of sharing exertion information with the leader[27], I was able to provide means to fill the knowledge gap between "what I think I know of my peers" and "what my peers are feeling at the moment." Sharing led to the group being more aware of everyone's exertion levels which filled the context for more informed paced decisions by leaders resulting in fewer drops while riding at stronger paces simultaneously.

It is essential to consider how to deliver information to riders and warrants further exploration. While both Paceguide and RPEView were effective, they provide different levels of information and require different amounts of "look-down" time to read the group effort. The former is a crucial consideration as riders who are interacting with their computers are not looking to the road (and riders) ahead. While there is no research regarding safe time windows for look-down touch screen interaction times for cyclists, car research suggests no longer than 2 seconds (Klauer et al., 2006; Zwahlen et al., 1988).

Scalability is another concern, especially for RPE View; as the group size increases, riders might have to start scrolling to see everyone in their group as the screen space is limited. RPE View had the highest self-reported group awareness scores; however, the future design must consider ways to avoid scrolling by prioritizing riders struggling or displaying the effort information in an order that mimics the current order of riders in the paceline. An initial approach to this problem is depicted in Figure 43. This design was generated from the feedback received in the discussion periods after the last study.

[27] And with everyone else in the group for my study

Lastly, recording and displaying exertion trends, similar to Star's Boundary Objects (Star & Griesemer, 1989), might be beneficial for riders too. As a rider moves from a high exertion role in the lead and moves to the wheel, their exertion levels will not immediately taper down to an active recovery state. The new leader could decrease the pace to accommodate the high exertion numbers on the wheel without the need for such action. A trend indication supports the natural decrease in exertion over time when an athlete moves from high intensity to active recovery closing the knowledge gap further.

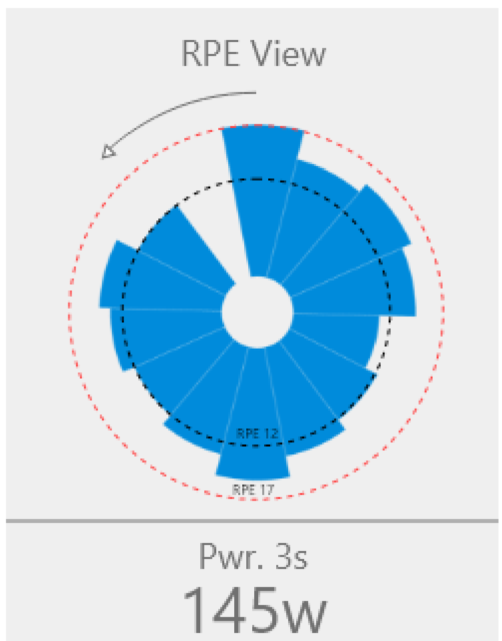

Figure 43 – Circular approach to RPE View addressing scalability (here showing 13 riders). Riders' RPEs are shown as a modified Rose Diagram. The leader's RPE is always at zero degrees. As riders rotate to the paceline their RPEs will move counter clockwise mimicking the rotation action in the real world. The gap divides the wheel and the leader roles, and the dash lines provide quick at a glance visualizations of any rider that is above or below RPEs 12 and 17.

However, different than Star's Boundary Objects or even Kolovison's (2020) sport specific Extraction Objects, the cycling work, did not focus on providing means for people with different roles to understand each other. Because riders rotate across all roles everyone understands the responsibilities and data needs of each position. Instead, I focused on creating means for an extra layer of communication to support sharing of information (exertion) among riders – a unique object shared across all riders – thus supporting their work as a team.

This added layer of exertion information transmitted through the bicycle computers creates a form of Workspace Awareness (Gutwin & Greenberg, 1996) that not only serves as another information dimension for the leader but also to create a shared understanding of the exertion levels of all the riders in the group.

Findings from the last study show that it is possible to make the knowledge gap of athletes smaller when their needs are evaluated in context, and design choices are made with their tasks in mind. Outcomes are valid across different group cycling disciplines such as Team Time Trial (TTT), Randonneuring and the grand tours such as Giro d'Italia and Tour de France. Other team sports that support in-loco data visualization, such as team rowing, could also benefit from exertion sharing however, the information delivery method should be redesigned to fit their work environment.

7.1 An Organizing Matrix of the Thesis Studies

Figure 44's design space serves as an organizing construct to better explain the difference between various components of the studies presented in my thesis. When comparing Aggregated vs. Raw presentation of information results indicated that both approaches were successful in supporting collaboration; however, raw shared content led to an increased cognitive load on participants either because they had to pay more attention to the content they were sharing (e.g., the live view in Over-the-Shoulder) or due to a larger quantity of exertion data in RPE View (although there were strategic advantages of knowing the exertion levels of all member individually). From the WYSIWIS vs. WYSINWIS lens, the first was able to *naturally* provide means to support sharing and collaboration among participants. However, it was not necessarily required, as seen by the Semantic Linking condition in the museum (and past work (Stefik et al., 1987)). Furthermore, WYSINWIS adds the capability of role-specific visualization (e.g., ship trajectories and metadata views in the maritime domain) and asynchronous interaction observed in the "look-ahead" behaviour in the SL condition in the museum.

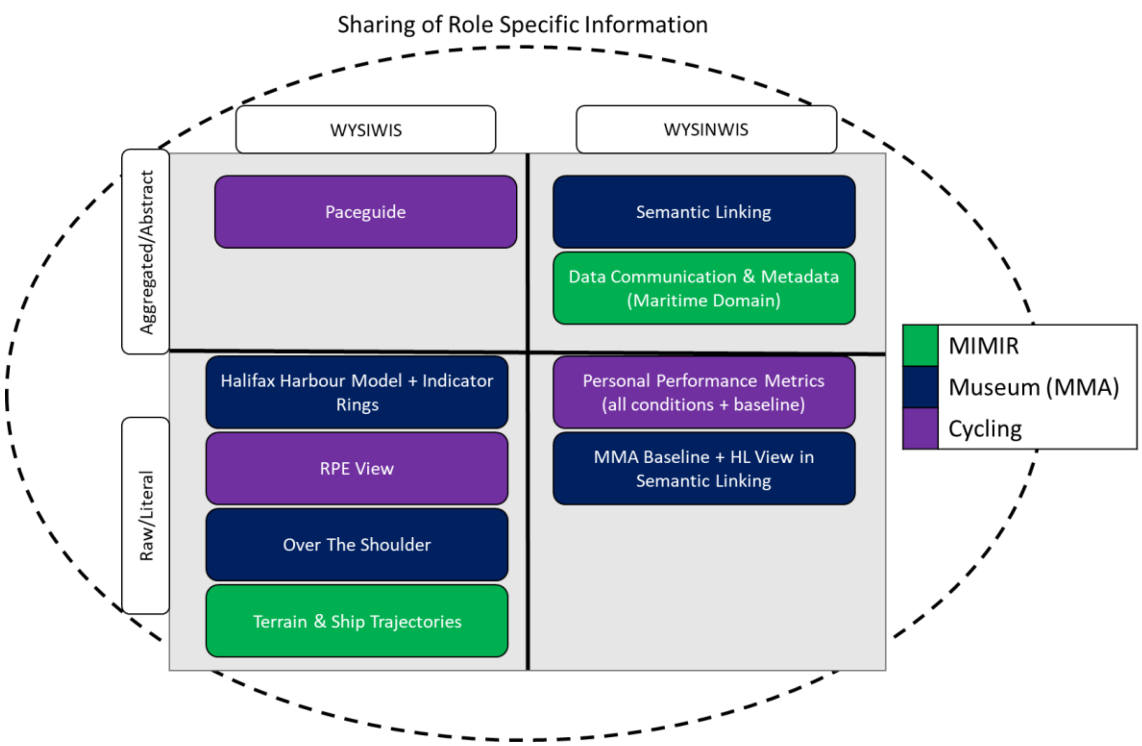

Figure 44 – Design space framing of my work which aims to support in role specific information sharing to users and enhance collaboration among them. Through my preliminary studies (MIMIR) and the core studies (MMA and Cycling), I explored interaction and sharing using both aggregate and raw data visualization in both WYSIWIS and WYSINWIS approaches.

Roles are fundamental to the success of CSCW systems (Guzdial et al., 2000; Zhu, 2006). It is the overarching concept on Figure 46 spanning across WYSIWIS and WYSINWIS approaches, and on the Raw and Aggregated representations. One of the foundations for the success of the Maritime Museum and Cycling studies was their intuitive role support. In both, roles are well defined; in MMA roles are established at the start of the experience (HL vs. no-HL), remain fixed[28] during the experiment, and require no training from the users. While in Road Cycling, roles are defined a priori, and users can and will freely transition between them throughout their collaborative work. In both approaches, roles were well-defined, promoting collaboration (Becht et al., 1999). Furthermore, the extra collaboration support when sharing exertion metrics led to a more performant team, another success from a CSCW and sports training perspective.

7.2 The Road Ahead

Future work should consider applications beyond the two domains explored in the three studies. Work should explore how the sharing methods created in the first study scale when there is more than one augmented reality headset present. Until now a one-to-many approach makes it simple to control which and how the virtual content is shared to non-AR users. Challenges arise when we have two or more headsets. Questions such as who controls the shared display, how to relinquish control to other AR devices, and how to indicate who is sharing content with non-AR users must then be addressed. Furthermore, another question that warrants exploration is whether non-AR users can 'detach' from whoever is sharing the AR content and promote individual exploration while maintaining

[28] As noted, outside a controlled study I would expect visitors to share the HMD by taking turns, especially in the baseline condition

engagement with their peers. Questions on how to indicate to non-AR users who is controlling the display and how to transition between WYSIWIS and WYSINWIS must be explored. Further work on merging the views of both, or many, AR views into a additive WYSIWIS view should also be considered.

Findings from the last study should also be explored outside the cycling environment. For example, sensing and sharing exertion in small displays, such as smart watches or heads-up AR displays, could benefit collocated teams working in hazardous environments that are physically demanding (Bhattarai et al., 2020). Noteworthy examples include first responders (e.g., fire crews), military, and search and rescue teams (B. Jones et al., 2020).

In the last study all participants had access to the shared information (WYSIWIS), future work should explore the impacts of showing the shared information only to the decision maker (leader) of the group (WYSINWIS). What are the impacts on performance, organization, and social aspects of this change in access?

7.2.1 *Intelligent* Exertion

A limitation of my last study was the manual input of the RPE values by the participants. While this works in research settings where the researcher constantly asks for the riders' exertion value, manual entry can be spotty at best. One option for new research is to explore if riders can retain this new effort information for future rides, so they do not need to use RPEView anymore as there is extra mental taxation to use it in the first place.

A more exciting path is for future work to explore ways to automate (partially or fully) the exertion sensing process. An interesting starting point is to explore other power-based metrics, such as Skiba's W'balance[29] , which could work as a starting point to judge if riders are resting or not through the changes in the W'balance value over time. Power is a critical metric in cycling, but novices do not widely use it because of the high cost of sensors (power meters). Access to current-generation power meters might be restricted due to high costs; research indicates (Araújo & Balbinot, 2022) that different approaches to sensing could lead to lower unit costs and broader usage of power sensing among enthusiasts and recreational road cyclists.

Power-based metrics alone might not be sufficient, however. Power can only measure the amount of energy expended during the ride and does not consider the rider's state before the rider starts (something RPE can take into account). There could be ways to retrieve this information using a combination (sensor fusion) of power metrics and heart rate data, such as heart rate variability metrics that gauge how tired a person is throughout the day (*Firstbeat Analytics*, n.d.).

Another interesting path is to use machine learning algorithms to model riders' exertion using power, heart rate, RPE, and other metrics throughout rides. This information can then be used to model the rider's exertion in future riders automating the manual input of implicit exertion values such as RPE by predicting the rider's current effort and sharing it with others. Riders can fine-tune the model as they ride by overriding predictions, thus leading to refined predictions in future rides.

[29] How much work above Critical Power a rider "has left". W'balance can be recharged by long periods of riding below CP.

7.2.2 Going Electric

Breaking away from group rides, the work presented here could also be integrated into e-bicycles for health promotion or socialization. Electric bicycles are an excellent tool for commuters or unfit people as they assist riders with their electric motors, thus reducing (or nulling) the effort required to go from point A to B.

By sensing the rider's effort and sharing it with the bicycle's motor controller, it could be possible to control the amount of pedal-assist or regenerative breaking transferred to the rider to target a specific exertion goal. Routes or training routines previously inaccessible for unfit riders can now be unlocked with partial electric support, possibly leading to more motivated riders.

Future research should explore using exertion-based electric pedal assist to bring together pairs or groups of riders with different fitness levels. Riders who are less fit than others could have just enough assistance to keep together with the group while simultaneously achieving training impact. More experienced riders would have their bikes with regen engaged, increasing the load on their pedals when they are not leading, or the pace is slower than they would typically go.

7.2.3 Ad-hoc Rides

The work in Chapter 6 focuses on a group of riders who know each other and have experience riding together. Riders within the group already have a pre-defined mental model of their peers' fitness abilities, albeit incomplete or not perfect, as seen by the studies in Chapters 5 and 6.

Future work should consider exploring the impacts of sharing metrics in groups unfamiliar with one another. Ad-hoc rides where riders meet up on a pre-determined start location during a given day are an interesting space for exploration, which was considered for my thesis however, riders were not happening regularly due to COVID-19.

CHAPTER 8 CONCLUSION

Throughout my thesis, I explored how sharing role specific information can fill the situational knowledge gap and ground collaboration in collocated individuals with different access to the same data. In the study presented in Chapter 3, I have demonstrated that museums with interactive augmented reality exhibits can use asymmetrical sharing to promote engagement between a patron wearing a head-mounted display and other members of their group. This is an important finding as it shows that museums do not necessarily need many headsets to promote these novel exhibits, which could be prohibitive with their budgets. Patrons in the study engaged with their peers while only one had access to the HMD while others could see proxy information on a wall-mounted display and projected it on the physical exhibit piece.

After, I switched domains from museums to group road cycling while maintaining the focus on shared experiences. In Chapter 5, I present the results of a Contextual Design study on the work done by group road cyclists while exercising together. Results are grouped using the five work models defined by the CD methodology (available on Appendices E through I). The study depicts a situation where riders can access sensor data while riding. However, due to the lack of sharing mechanisms, the information is contained within each person, creating silos of knowledge. I have identified future research and design challenges to support riders in areas such as exertion sensing and sharing, communication, data presentation, role-based tasks and needs, and pacing support.

The last study in my thesis (Chapter 6) focuses on one of the issues identified in the Contextual Design study: the lack of exertion sensing and sharing among riders leads to attrition between riders and sub-optimized paces. In this study, I started from the principle that we can and should leverage the current ubiquitous hardware instead of creating new lab-only, impractical solutions and propose two ways of sharing effort among the peloton designed with their computers in mind.

Results indicated that sensing and sharing effort leads to more smooth riders with fewer drops and increased awareness of everyone's exertion levels. Interestingly, results also show that riders with sharing were able to have higher normalized power than without. This difference indicates that the extra information allowed riders to push the power envelope further without sacrificing the group's cohesion. Outcomes from this study should be directly applicable to different group cycling modalities such as TTT and Randonneuring and other group sports that support in-situ data visualizations.

Future work can address the challenges faced when scaling the AR sharing techniques to two or more headsets, the other issues identified in the Contextual Design study, automated ways of sensing implicit exertion values such as body sensations and bringing exertion values to other group sports or collocated teams such as search and rescue crews.

BIBLIOGRAPHY

Ackerman, M. S. (2000). Intellectual challenge of CSCW: the gap between social requirements and technical feasibility. *Human-Computer Interaction*, *15*(2–3), 179–203. https://doi.org/10.1207/S15327051HCI1523_5

Adnan, M., Just, M., & Baillie, L. (2016). Investigating Time Series Visualisations to Improve the User Experience. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*, 5444–5455. https://doi.org/10.1145/2858036.2858300

Alberton, C. L., Antunes, A. H., Beilke, D. D., Pinto, S. S., Kanitz, A. C., Tartaruga, M. P., & Martins Kruel, L. F. (2013). Maximal and Ventilatory Thresholds of Oxygen Uptake and Rating of Perceived Exertion Responses to Water Aerobic Exercises. *Journal of Strength and Conditioning Research*, *27*(7), 1897–1903. https://doi.org/10.1519/JSC.0b013e3182736e47

Allen, H., Coggan, A. R., & McGregor, S. (2019). *Training and racing with a power meter*. VeloPress.

Alvermann, J. (2016). Mobile Media in the Museum Space: The Example of the Neanderthal Museum's App "Neanderthal+." *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*, 1509–1512. https://doi.org/10.1145/2968219.2974048

Aoki, P. M., Grinter, R. E., Hurst, A., Szymanski, M. H., Thornton, J. D., & Woodruff, A. (2002). Sotto Voce: Exploring the interplay of Conversation and Mobile Audio Spaces. *SIGCHI Conference on Human Factors in Computing Systems: Changing Our World, Changing Ourselves*, *1*, 431–438. https://doi.org/10.1145/503376.503454

Aranki, D., Balakrishnan, U., Sarver, H., Serven, L., Asuncion, C., Du, K., Gruis, C., Peh, G. X., Xiao, Y., & Bajcsy, R. (2017). RunningCoach. *Proceedings of the 11th EAI International Conference on Pervasive Computing Technologies for Healthcare*, 325–334. https://doi.org/10.1145/3154862.3154935

Araújo, M. O., & Balbinot, A. (2022). *Proposal of a Low Profile Piezoelectric Based Insole to Measure the Perpendicular Force Applied by a Cyclist* (pp. 813–819). Springer, Cham. https://doi.org/10.1007/978-3-030-70601-2_123

Arroyo Ernesto and Righi, V. and T. R. and B. J. (2011). A Remote Multi-touch Experience to Support Collaboration between Remote Museum Visitors. In N. and J. J. and N. N. and P. P. and W. M. Campos Pedro and Graham (Ed.), *Human-Computer Interaction – INTERACT 2011* (pp. 462–465). Springer Berlin Heidelberg.

At Home Cycling & Running Virtual Training & Workout Game App. (n.d.). Retrieved May 1, 2022, from https://www.zwift.com/ca

Azuma, R. T. (1997). A Survey of Augmented Reality. *Presence: Teleoperators and Virtual Environments*, *6*(4), 355–385. https://doi.org/10.1162/pres.1997.6.4.355

Baca, A., Dabnichki, P., Heller, M., & Kornfeind, P. (2009). Ubiquitous computing in sports: A review and analysis. *Journal of Sports Sciences*, *27*(12), 1335–1346. https://doi.org/10.1080/02640410903277427

Becht, M., Gurzki, T., Klarmann, J., & Muscholl, M. (1999). ROPE: role oriented programming environment for multiagent systems. *Proceedings Fourth IFCIS International Conference on Cooperative Information Systems. CoopIS 99 (Cat. No.PR00384)*, 325–333. https://doi.org/10.1109/COOPIS.1999.792181

Beck, F., Burch, M., & Weiskopf, D. (2016). A matrix-based visual comparison of time series sports data. *VMV 2016 - Vision, Modeling and Visualization*, 53–60. https://doi.org/10.2312/vmv.20161342

Beecham, R., & Wood, J. (2014). Characterising group-cycling journeys using interactive graphics. *Transportation Research Part C: Emerging Technologies*, *47*(P2), 194–206. https://doi.org/10.1016/j.trc.2014.03.007

Beecham, R., Wood, J., & Bowerman, A. (2012). A visual analytics approach to understanding cycling behaviour. *IEEE Conference on Visual Analytics Science and Technology 2012, VAST 2012 - Proceedings*, 207–208. https://doi.org/10.1109/VAST.2012.6400550

Beecken, J., & Reinhardt, A. (2019). How Smooth is my Ride? Detecting Bikeway Conditions from Smartphone Video Streams. *2019 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, 65–70. https://doi.org/10.1109/PERCOMW.2019.8730869

Benko, H., Ishak, E. W., & Feiner, S. (2004). Collaborative mixed reality visualization of an archaeological excavation. *ISMAR 2004: Proceedings of the Third IEEE and ACM International Symposium on Mixed and Augmented Reality*, 132–140. https://doi.org/10.1109/ISMAR.2004.23

Beyer, H., & Holtzblatt, K. (1997). *Contextual Design: Defining Customer-Centered Systems*. Morgan Kaufmann Publishers Inc.

Bhattarai, M., Jensen-Curtis, A. R., & Martinez-Ramon, M. (2020). An embedded deep learning system for augmented reality in firefighting applications. *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 1224–1230. https://doi.org/10.1109/ICMLA51294.2020.00193

Billinghurst, M., & Kato, H. (1999). Collaborative Mixed Reality. *Proceedings of the First International Symposium on Mixed Reality*, 261–284. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.107.682&rep=rep1&type=pdf

Bødker, S. (2006). When second wave HCI meets third wave challenges. *ACM International Conference Proceeding Series*, *189*, 1–8. https://doi.org/10.1145/1182475.1182476

BORG, G. A. V. (1982). Psychophysical bases of perceived exertion. *Medicine & Science in Sports & Exercise*, *14*(5), 377???381. https://doi.org/10.1249/00005768-198205000-00012

Borszcz, F., Tramontin, A., Bossi, A., Carminatti, L., & Costa, V. (2018). Functional Threshold Power in Cyclists: Validity of the Concept and Physiological Responses. *International Journal of Sports Medicine*, *39*(10), 737–742. https://doi.org/10.1055/s-0044-101546

Braunhofer, M., Ricci, F., Lamche, B., & Wörndl, W. (2015). A Context-Aware Model for Proactive Recommender Systems in the Tourism Domain. *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct - MobileHCI '15*, 1070–1075. https://doi.org/10.1145/2786567.2794332

Brooke, J. (1996). SUS - A quick and dirty usability scale. *Usability Evaluation in Industry*, *189*(194), 4–7. https://doi.org/10.1002/hbm.20701

Brown, P. J. (1996). *The Stick-e Document: a Framework for Creating Context-aware Applications*. https://www.semanticscholar.org/paper/The-Stick-e-Document%3A-a-Framework-for-Creating-Brown/b66dd0764d2a0135e977804b3345b242cf1ee3bf

Burrell, J., & Gay, G. K. (2002). E-graffiti: evaluating real-world use of a context-aware system. *Interacting with Computers*, *14*(4), 301–312. https://doi.org/10.1016/S0953-5438(02)00010-3

Cespedes, S., Salamanca, J., Yanez, A., & Vinasco, D. (2019). Group cycling meets technology: A cooperative cycling cyber-physical system. *IEEE Transactions on Intelligent Transportation Systems*, *20*(8), 3178–3188. https://doi.org/10.1109/TITS.2018.2874394

Chalmers, M., & Galani, A. (2004). Seamful interweaving: heterogeneity in the theory and design of interactive systems. *Proc. DIS 2004*, *August*, 243–252. https://doi.org/http://doi.acm.org/10.1145/1013115.1013149

Chandler, T., Cordeil, M., Czauderna, T., Dwyer, T., Glowacki, J., Goncu, C., Klapperstueck, M., Klein, K., Marriott, K., Schreiber, F., & Wilson, E. (2015). Immersive Analytics. *2015 Big Data Visual Analytics, BDVA 2015*. https://doi.org/10.1109/BDVA.2015.7314296

Cheok, A. D., Wang Weihua, Xubo Yang, Prince, S., Fong Siew Wan, Billinghurst, M., & Kato, H. (n.d.). Interactive theatre experience in embodied + wearable mixed reality space. *Proceedings. International Symposium on Mixed and Augmented Reality*, 59–317. https://doi.org/10.1109/ISMAR.2002.1115073

Cheverst, K., Davies, N., Mitchell, K., Friday, A., & Efstratiou, C. (2000). Developing a Context-Aware Electronic Tourist Guide: Some Issues and Experiences. *SIGCHI Conference on Human Factors in Computing Systems*, *2*(1), 17–24. https://doi.org/10.1145/332040.332047

Circumstances surrounding cycling fatalities in Canada, 2006 to 2017. (n.d.). Retrieved May 25, 2022, from https://www150.statcan.gc.ca/n1/pub/82-625-x/2019001/article/00009-eng.htm

Colley, A., Woźniak, P. W., Kiss, F., & Häkkilä, J. (2018). Shoe integrated displays. *Proceedings of the 10th Nordic Conference on Human-Computer Interaction - NordiCHI '18*, 39–46. https://doi.org/10.1145/3240167.3240216

Čopič Pucihar, K., & Kljun, M. (2018). ART for Art: Augmented Reality Taxonomy for Art and Cultural Heritage. *Augmented Reality Art*, *February*, 73–94. https://doi.org/10.1007/978-3-319-69932-5_3

Crowell, H. P., Milnert, C. E., Hamill, J., & Davis, I. S. (2010). Reducing impact loading during running with the use of real-time visual feedback. *Journal of Orthopaedic and Sports Physical Therapy*, *40*(4), 206–213. https://doi.org/10.2519/jospt.2010.3166

Dancu, A., Vechev, V., Ayça Ünlüer, A., Nilson, S., Nygren, O., Eliasson, S., Barjonet, J.-E. E., Marshall, J., Fjeld, M., Ünlüer, A. A., Nilson, S., Nygren, O., Eliasson, S., Barjonet, J.-E. E., Marshall, J., Fjeld, M., Ayça Ünlüer, A., Nilson, S., Nygren, O., … Fjeld, M. (2015). Gesture Bike: Examining Projection Surfaces and Turn Signal Systems for Urban Cycling. *Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces - ITS '15*, 151–159. https://doi.org/10.1145/2817721.2817748

Dancu, A., Vechev, V., Ünlüer, A. A., Nilson, S., Nygren, O., Eliasson, S., Barjonet, J.-E., Marshall, J., & Fjeld, M. (2015). Gesture Bike. *Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces - ITS '15*, 151–159. https://doi.org/10.1145/2817721.2817748

Davidson, K., & McNaughton, L. (2000). Deep Water Running Training and Road Running Training Improve Vo2max in Untrained Women. *Journal of Strength and Conditioning Research*. https://doi.org/10.1519/00124278-200005000-00012

de Brouwer, M., Ongenae, F., Daneels, G., Municio, E., Famaey, J., Latré, S., & de Turck, F. (2018). *Personalized Real-Time Monitoring of Amateur Cyclists on Low-End Devices*. 1833–1840. https://doi.org/10.1145/3184558.3191648

Debord, G. (1955). Introduction to a Critique of Urban Geography. In *Critical Geographies: A Collection of Readings*. https://doi.org/10.1109/ICUAS.2016.7502677

Dey, A. K. (1998). Context-Aware Computing: The CyberDesk Project. *AAAI 1998*. https://www.semanticscholar.org/paper/Context-Aware-Computing%3A-The-CyberDesk-Project-Dey/9d49c8564784197ebbe7df6395544e54f691cbae

Dey, A. K. (1999). *Providing Architectural Support for Building Context-Aware Applications*.

Dey, A. K. (2001). Understanding and Using Context. *Personal and Ubiquitous Computing*, *5*(1), 4–7. https://doi.org/10.1007/s007790170019

Doering, Z. (2004). Results of the 2004 Smithsonian-wide Survey of Museum Visitors. *Smithsonian Institution, Office of Policy and Analysis, Washington, DC*, 4–5.

Dourish, P. (2004). What we talk about when we talk about context. *Personal and Ubiquitous Computing*, *8*(1), 19–30. https://doi.org/10.1007/s00779-003-0253-8

Dourish, P. (2006). *Implications for Design*.

Dustdar, M., & Rosenberg, S. (2007). A survey on context-aware systems. In *Int. J. Ad Hoc and Ubiquitous Computing* (Vol. 2, Issue 4). http://www.infosys.tuwien.ac.at/Staff/rosenberg.

Eisenman, S. B., Miluzzo, E., Lane, N. D., Peterson, R. A., Ahn, G.-S., & Campbell, A. T. (2007). The BikeNet mobile sensing system for cyclist experience mapping. *Proceedings of the 5th International Conference on Embedded Networked Sensor Systems - SenSys '07*, *6*(1), 87. https://doi.org/10.1145/1322263.1322273

Eisenman, S. B., Miluzzo, E., Lane, N. D., Peterson, R. A., Ahn, G.-S., & Campbell, A. T. (2009). BikeNet. *ACM Transactions on Sensor Networks*, *6*(1), 1–39. https://doi.org/10.1145/1653760.1653766

el Esawey, M., Lim, C., & Sayed, T. (2015). Development of a cycling data model: City of Vancouver case study. *Canadian Journal of Civil Engineering*, *42*(12), 1000–1010. https://doi.org/10.1139/cjce-2015-0065

Firstbeat Analytics. (n.d.). Retrieved June 26, 2022, from https://www.firstbeatanalytics.com/en/

Fisken, A. L., Waters, D. L., Hing, W. A., & Keogh, J. W. (2016). Perceptions towards aqua-based exercise among older adults with osteoarthritis who have discontinued participation in this exercise mode. *Australasian Journal on Ageing*, *35*(1), 12–17. https://doi.org/10.1111/ajag.12167

Fister, I., & Fister Jr, I. (2012). Concept of drafting detection system in Ironmans. *ArXiv Preprint ArXiv:1206.1851*.

Fleck, M., Frid, M., Kindberg, T., O'Brien-Strain, E., Rajani, R., & Spasojevic, M. (2002). From informing to remembering: ubiquitous systems in interactive museums. *IEEE Pervasive Computing*, *1*(2), 13–21. https://doi.org/10.1109/MPRV.2002.1012333

Fosh, L., Benford, S., & Koleva, B. (2016). *Supporting Group Coherence in a Museum Visit*. 1–12. https://doi.org/10.1145/2818048.2819970

Franz, J., Alnusayri, M., Malloch, J., & Reilly, D. (2019). A Comparative Evaluation of Techniques for Sharing AR Experiences in Museums. *Proceedings of the ACM on Human-Computer Interaction*, *3*(CSCW), 1–20. https://doi.org/10.1145/3359226

Franz, J. M., Malloch, J. W., Reilly, D., & Sanches, V. (2018). SpaceHopper. *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–4. https://doi.org/10.1145/3170427.3186509

Franz, J., Malloch, J., Nedel, L., & Reilly, D. (2017). More Than Blips on the Radar: Exploring Immersive Visualization for Maritime Decision Making. *Electronic Proceedings of Immersive Analytics Workshop at VIS*.

Franz, J., Malloch, J., & Reilly, D. (2017). Exploring Shared Immersive Visualization in AR. *Extended Abstracts of VIS 2017*.

Franz, J., Malloch, J., & Reilly, D. (2020). Compensating for Perspective-based Distortion on Large Interactive Floor Displays: The SpaceHopper Field Experiment. *IMX 2020 - Proceedings of the 2020 ACM International Conference on Interactive Media Experiences*, 45–54. https://doi.org/10.1145/3391614.3393651

Franz, J., & Reilly, D. (2022). Ride With Me: Exploring Group Road Cycling Through Contextual Design. *Designing Interactive Systems Conference*, 1612–1625. https://doi.org/10.1145/3532106.3533476

Galani, A., & Chalmers, M. (2013). Empowering the remote visitor: supporting social museum experiences among local and remote visitors. In *Museums in a digital age* (pp. 172–182). Routledge.

Geller, T. (2006). Interactive Tabletop Exhibits in Museums and Galleries. *IEEE Computer Graphics and Applications*, *26*(5), 6–11. https://doi.org/10.1109/MCG.2006.111

Gerval, J.-P., & le Ru, Y. (2015). Fusion of Multimedia and Mobile Technology in Audioguides for Museums and Exhibitions. In D. Sharma, M. Favorskaya, L. C. Jain, & R. J. Howlett (Eds.), *Fusion of Smart, Multimedia and Computer Gaming Technologies: Research, Systems and Perspectives* (pp. 173–205). Springer International Publishing. https://doi.org/10.1007/978-3-319-14645-4_8

Giro d'Italia 2022 | Official site. (n.d.). Retrieved May 11, 2022, from https://www.giroditalia.it/en/

GoldenCheetah. (n.d.). Retrieved May 1, 2022, from https://www.goldencheetah.org/

Grinter, R. E., Aoki, P. M., Szymanski, M. H., Thornton, J. D., Woodruff, A., & Hurst, A. (2002). Revisiting the visit: understanding how technology can shape the museum visit. *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work - CSCW '02*, *Proceeding*(1), 146. https://doi.org/10.1145/587078.587100

Gudmundsson, J., & Horton, M. (2017). Spatio-Temporal Analysis of Team Sports. *ACM Computing Surveys*, *50*(2), 1–34. https://doi.org/10.1145/3054132

Gutwin, C., & Greenberg, S. (1996). Workspace awareness for groupware. *Conference on Human Factors in Computing Systems - Proceedings*, 208–209. https://doi.org/10.1145/257089.257284

Guzdial, M., Rick, J., & Kerimbaev, B. (2000). Recognizing and supporting roles in CSCW. *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work - CSCW '00*, 261–268. https://doi.org/10.1145/358916.358997

Häuslschmid, R., Fritzsche, B., & Butz, A. (2018). Can a Helmet-mounted Display Make Motorcycling Safer? *Proceedings of the 2018 Conference on Human Information Interaction&Retrieval - IUI '18*, 467–476. https://doi.org/10.1145/3172944.3172963

Hensel, K. A. (1987). *Families in a museum: Interactions and conversations at displays*. Teachers College, Columbia University.

Herbst, I., Braun, A.-K., McCall, R., & Broll, W. (2008). TimeWarp. *Proceedings of the 10th International Conference on Human Computer Interaction with Mobile Devices and Services - MobileHCI '08*, 235. https://doi.org/10.1145/1409240.1409266

Hernández–Melgarejo, G., Flores–Hernández, D. A., Luviano–Juárez, A., Castañeda, L. A., Chairez, I., & di Gennaro, S. (2020). Mechatronic design and implementation of a bicycle virtual reality system. *ISA Transactions*, *97*, 336–351. https://doi.org/10.1016/j.isatra.2019.08.002

Hindmarsh, J., Heath, C., vom Lehn, D., & Cleverly, J. (2002). Creating assemblies: aboard the Ghost Ship. *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work - CSCW '02*, 156. https://doi.org/10.1145/587078.587101

Hood, M. (1983). Staying Away: Why People Choose Not to Visit Museums | Mendeley. *Museum News*, *61*(4), 50–57. https://contentstore.cla.co.uk/secure/link?id=75ddf2cd-d20c-e811-80cd-005056af4099

Hornecker, E., & Nicol, E. (2012). What do lab-based user studies tell us about in-the-wild behavior? *Designing Interactive Systems Conference*, 358–367. https://doi.org/10.1145/2317956.2318010

Horrey, W. J., & Wickens, C. D. (2004). Focal and Ambient Visual Contributions and Driver Visual Scanning in Lane Keeping and Hazard Detection. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *48*(19), 2325–2329. https://doi.org/10.1177/154193120404801926

Huang, D., Tory, M., Adriel Aseniero, B., Bartram, L., Bateman, S., Carpendale, S., Tang, A., & Woodbury, R. (2015). Personal visualization and personal visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, *21*(3), 420–433. https://doi.org/10.1109/TVCG.2014.2359887

Irwin, B. C., Scorniaenchi, J., Kerr, N. L., Eisenmann, J. C., & Feltz, D. L. (2012). Aerobic Exercise Is Promoted when Individual Performance Affects the Group: A Test of the Kohler Motivation Gain Effect. *Annals of Behavioral Medicine*, *44*(2), 151–159. https://doi.org/10.1007/s12160-012-9367-4

Ito, K., Nishimura, H., & Ogi, T. (2015). Head-up display for motorcycle navigation. *SIGGRAPH Asia 2015 Head-Up Displays and Their Applications on - SA '15*, 1–3. https://doi.org/10.1145/2818406.2818415

Janetzko, H., Sacha, D., Stein, M., Schreck, T., Keim, D. A., & Deussen, O. (2015). Feature-driven visual analytics of soccer data. *2014 IEEE Conference on Visual Analytics Science and Technology, VAST 2014 - Proceedings*, 13–22. https://doi.org/10.1109/VAST.2014.7042477

Jeon, W., & Rajamani, R. (2018). Rear vehicle tracking on a bicycle using active sensor orientation control. *IEEE Transactions on Intelligent Transportation Systems*, *19*(8), 2638–2649. https://doi.org/10.1109/TITS.2017.2764006

Johansen, R. (1988). *GroupWare: Computer Support for Business Teams*. The Free Press.

Jones, B., Tang, A., Neustaedter, C., Antle, A. N., & McLaren, E.-S. (2020). *Designing Technology for Shared Communication and Awareness in Wilderness Search and Rescue* (pp. 175–194). https://doi.org/10.1007/978-3-030-45289-6_9

Jones, E. M., Selker, T., & Chung, H. (2007). What you said about where you shook your head. *CHI '07 Extended Abstracts on Human Factors in Computing Systems - CHI '07*, 2477. https://doi.org/10.1145/1240866.1241027

Jose, R., Lee, G. A., & Billinghurst, M. (2016). A comparative study of simulated augmented reality displays for vehicle navigation. *Proceedings of the 28th Australian Conference on Computer-Human Interaction - OzCHI '16*, 40–48. https://doi.org/10.1145/3010915.3010918

Kadomura, A., Ichioka, Y., Tsukada, K., Rekimoto, J., & Siio, I. (2014). *AwareCycle: Application for Sports Visualization Using an Afterimage Display Attached to the Wheel of a Bicycle* (pp. 256–264). Springer, Cham. https://doi.org/10.1007/978-3-319-07227-2_25

Katifori, A., Perry, S., Vayanou, M., Pujol, L., Chrysan, A., Kourtis, V., & Ioannidis, Y. (2016). Cultivating Mobile-Mediated Social Interaction in the Museum: Towards Group-Based Digital Storytelling Experiences. *Annual Conference of Museums and the Web: MW2016*. http://eprints.gla.ac.uk/143241/

Kerr, N. L., & Hertel, G. (2011). The Köhler Group Motivation Gain: How to Motivate the 'Weak Links' in a Group. *Social and Personality Psychology Compass*, *5*(1), 43–55. https://doi.org/10.1111/j.1751-9004.2010.00333.x

Khot, R. A., Hjorth, L., & Mueller, F. (2020). Shelfie. *ACM Transactions on Computer-Human Interaction*, *27*(3), 1–52. https://doi.org/10.1145/3379539

Kiss, F., Dominiak, J., Scheerer, F., & Schmidt, A. (2019). Clairbuoyance : Improving Directional Perception for Swimmers. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, 1–12.

Klauer, C., Dingus, T. A., Neale, V. L., Sudweeks, J. D., Ramsey, D. J., & others. (2006). *The impact of driver inattention on near-crash/crash risk: An analysis using the 100-car naturalistic driving study data*.

Köhler, O. (1927). Über den Gruppenwirkungsgrad der menschlichen Körperarbeit und die Bedingung optimaler Kollektivkraftreaktion. *Industrielle Psychotechnik*.

Kolovson, S., Liang, C., Munson, S. A., & Starbird, K. (2020). Personal data and power asymmetries in U.S. Collegiate sports teams. *Proceedings of the ACM on Human-Computer Interaction*, *4*(GROUP). https://doi.org/10.1145/3375185

Koutsabasis, P. (2017). Empirical Evaluations of Interactive Systems in Cultural Heritage. *International Journal of Computational Methods in Heritage Science (IJCMHS)*, *1*(1), 100–122. https://doi.org/10.4018/IJCMHS.2017010107

Lage, M., Ono, J. P., Cervone, D., Chiang, J., Dietrich, C., & Silva, C. T. (2016). StatCast Dashboard: Exploration of Spatiotemporal Baseball Data. *IEEE Computer Graphics and Applications*, *36*(5), 28–37. https://doi.org/10.1109/MCG.2016.101

Legg, P. A., Chung, D. H. S., Parry, M. L., Jones, M. W., Long, R., Griffiths, I. W., & Chen, M. (2012). MatchPad: Interactive glyph-based visualization for real-time sports performance analysis. *Computer Graphics Forum*, *31*(3 PART 4), 1255–1264. https://doi.org/10.1111/j.1467-8659.2012.03118.x

Liu, Y. C., & Wen, M. H. (2004). Comparison of head-up display (HUD) vs. head-down display (HDD): Driving performance of commercial vehicle operators in Taiwan. *International Journal of Human Computer Studies*, *61*(5), 679–697. https://doi.org/10.1016/j.ijhcs.2004.06.002

Live Segments – Strava Support. (n.d.). Retrieved May 11, 2022, from https://support.strava.com/hc/en-us/articles/207343830-Live-Segments

Lo, W. H., Zollmann, S., Regenbrecht, H., & Loos, M. (2019). From Lab to Field: Demonstrating Mixed Reality Prototypes for Augmented Sports Experiences. *The 17th International Conference on Virtual-Reality Continuum and Its Applications in Industry*, 1–2. https://doi.org/10.1145/3359997.3365728

López Sintas, J., García Álvarez, E., & Pérez Rubiales, E. (2014). Art museum visitors: interaction strategies for sharing experiences. *Museum Management and Curatorship*, *29*(3), 241–259. https://doi.org/10.1080/09647775.2014.919175

Losada, A. G., Theron, R., & Benito, A. (2016). BKViz: A Basketball Visual Analysis Tool. *IEEE Computer Graphics and Applications*, *36*(6), 58–68. https://doi.org/10.1109/MCG.2016.124

Luka, M. E., & Lilley, B. (2018). The NiS+TS Psychogeographer's Table: Countering the Official Halfiax Explosion Archive. *Public*, *57*(1), 295.

Macedo, A. Q., Marinho, L. B., & Santos, R. L. T. (2015). Context-Aware Event Recommendation in Event-based Social Networks. *Proceedings of the 9th ACM Conference on Recommender Systems - RecSys '15*, 123–130. https://doi.org/10.1145/2792838.2800187

Macintosh, B. R., Neptune, R. R., & Horton, J. F. (2000). Cadence, power, and muscle activation in cycle ergometry. *Medicine & Science in Sports & Exercise*, *32*(7), 1281–1287. https://doi.org/10.1097/00005768-200007000-00015

Madsen, J. B., & Madsen, C. B. (2015). Handheld Visual Representation of a Castle Chapel Ruin. *Journal on Computing and Cultural Heritage*. https://doi.org/10.1145/2822899

Matviienko, A., & Ali, A. el. (2019). *NaviBike : Comparing Unimodal Navigation Cues for Child Cyclists*. 1–12.

Matviienko, A., Ananthanarayan, S., Borojeni, S. S., Feld, Y., Heuten, W., & Boll, S. (2018). *Augmenting bicycles and helmets with multimodal warnings for children*. 1–13. https://doi.org/10.1145/3229434.3229479

McManus, P. (1989). Oh, Yes, They Do: How Museum Visitors Read Labels and Interact with Exhibit Texts. *Curator*, *32*(3), 15. https://doi.org/10.1111/j.2151-6952.1989.tb00718.x

McManus, P. M. (1987). It's the company you keep …: The social determination of learning-related behaviour in a science museum. *International Journal of Museum Management and Curatorship*, *6*(3), 263–270. https://doi.org/10.1080/09647778709515076

Medenica, Z., Kun, A. L., Paek, T., & Oskar. (2011). Augmented Reality vs. Street Views: A Driving Simulator Study. *MobileHCI*.

Mieras, M. E., Heesch, M. W. S., & Slivka, D. R. (2014). Physiological and Psychological Responses to Outdoor vs. Laboratory Cycling. *Journal of Strength and Conditioning Research*, *28*(8), 2324–2329. https://doi.org/10.1519/JSC.0000000000000384

Milgram, P., Takemura, H., Utsumi, a, & Kishino, F. (1994). Mixed Reality (MR) Reality-Virtuality (RV) Continuum. *Systems Research*, *2351*(Telemanipulator and Telepresence Technologies), 282–292. https://doi.org/10.1.1.83.6861

Municio, E., Daneels, G., de Brouwer, M., Ongenae, F., de Turck, F., Braem, B., Famaey, J., & Latre, S. (2019). Continuous Athlete Monitoring in Challenging Cycling Environments Using IoT Technologies. *IEEE Internet of Things Journal*, *6*(6), 10875–10887. https://doi.org/10.1109/JIOT.2019.2942761

Niederer, C., Rind, A., & Aigner, W. (2016). Multi-device visualisation design for climbing self-assessment. *Proceedings of the International Conference on Information Visualisation*, *2016-August*, 171–176. https://doi.org/10.1109/IV.2016.34

Nilsson, S., Johansson, B., & Jonsson, A. (2009). Using AR to support cross-organisational collaboration in dynamic tasks. *2009 8th IEEE International Symposium on Mixed and Augmented Reality*, 3–12. https://doi.org/10.1109/ISMAR.2009.5336522

Nissenbaum, H. (2004). Washington law review: Privacy as contextual integrity. *Washington Law Review*, *79*(1), 119–157. https://nyuscholars.nyu.edu/en/publications/privacy-as-contextual-integrity

Norman, G. (2010). Likert scales, levels of measurement and the “laws” of statistics. *Advances in Health Sciences Education*, *15*(5), 625–632. https://doi.org/10.1007/s10459-010-9222-y

Official website of Tour de France 2022. (n.d.). Retrieved May 11, 2022, from https://www.letour.fr/en/

Ogi, T. (2015). Design and evaluation of HUD for motorcycle using immersive simulator. *SIGGRAPH Asia 2015 Head-Up Displays and Their Applications on - SA ’15*, 1–2. https://doi.org/10.1145/2818406.2818411

Oral Kaplan, Goshiro Yamamoto, Yasuhide Yoshitake, Takafumi Taketomi, C. S. and H. K. (2016). *In-Situ Visualization of Pedaling Forces on Cycling Training Videos*.

Orlikowski, W. J., & Gash, D. C. (1994). Technological Frames: Making Sense of Information Technology in Organizations. *ACM Transactions on Information Systems (TOIS)*, *12*(2), 174–207. https://doi.org/10.1145/196734.196745

Packer, J., & Ballantyne, R. (2005). Solitary vs. Shared: Exploring the Social Dimension of Museum Learning. *Curator: The Museum Journal*, *48*(2), 177–192. https://doi.org/10.1111/j.2151-6952.2005.tb00165.x

Pascoe, J., Ryan, N., & Morse, D. (1999). *Issues in Developing Context-Aware Computing* (pp. 208–221). https://doi.org/10.1007/3-540-48157-5_20

Pascoe, J., Ryan, N., & Morse, D. (2000). Using While Moving: HCI Issues in Fieldwork Environments. *ACM Transactions on Computer-Human Interaction*, *7*(3), 417–437. https://doi.org/10.1145/355324.355329

Pauzié, A. (2008). A method to assess the driver mental workload: The driving activity load index (DALI). *IET Intelligent Transport Systems*, *2*(4), 315. https://doi.org/10.1049/iet-its:20080023

Pedersen, I., Gale, N., Mirza-Babaei, P., & Reid, S. (2017). More than Meets the Eye: The Benefits of Augmented Reality and Holographic Displays for Digital Cultural Heritage. *Journal on Computing and Cultural Heritage*, *10*(2), 1–15. https://doi.org/10.1145/3051480

Pejovic, V., & Musolesi, M. (2014). InterruptMe. *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '14 Adjunct*, 897–908. https://doi.org/10.1145/2632048.2632062

PEKARIK, A. J., DOERING, Z. D., & KARNS, D. A. (1999). Exploring Satisfying Experiences in Museums. *Curator: The Museum Journal*, *42*(2), 152–173. https://doi.org/10.1111/j.2151-6952.1999.tb01137.x

Pelowski, M., Liu, T., Palacios, V., & Akiba, F. (2014). When a body meets a body: An exploration of the negative impact of social interactions on museum experiences of art. *International Journal of Education and the Arts*, *15*, 1–48.

Perin, C., Vuillemot, R., & Fekete, J. D. (2013). SoccerStories: A kick-off for visual soccer analysis. *IEEE Transactions on Visualization and Computer Graphics*, *19*(12), 2506–2515. https://doi.org/10.1109/TVCG.2013.192

Perin, C., Vuillemot, R., Stolper, C. D., Stasko, J. T., Wood, J., & Carpendale, S. (2018). State of the Art of Sports Data Visualization. *Computer Graphics Forum*, *37*(3), 663–686. https://doi.org/10.1111/cgf.13447

Pierdicca, R., Bekele, M. K., Malinverni, E. S., Gain, J., & Frontoni, E. (2018). A Survey of Augmented, Virtual, and Mixed Reality for Cultural Heritage. *Journal on Computing and Cultural Heritage*. https://doi.org/10.1145/3145534

Pileggi, H., Stolper, C. D., Boyle, J. M., & Stasko, J. T. (2012). SnapShot: Visualization to propel ice hockey analytics. *IEEE Transactions on Visualization and Computer Graphics*, *18*(12), 2819–2828. https://doi.org/10.1109/TVCG.2012.263

Polk, T., Yang, J., Hu, Y., & Zhao, Y. (2014). TenniVis: Visualization for tennis match analysis. *IEEE Transactions on Visualization and Computer Graphics*, *20*(12), 2339–2348. https://doi.org/10.1109/TVCG.2014.2346445

Poppinga, B., Pielot, M., & Boll, S. (2009). *Tacticycle – A Tactile Display for Supporting Tourists on a Bicycle Trip*. ACM.

Probst, L., Rauschenbach, F., Schuldt, H., Seidenschwarz, P., & Rumo, M. (2018). Integrated Real-Time Data Stream Analysis and Sketch-Based Video Retrieval in Team Sports. *2018 IEEE International Conference on Big Data (Big Data)*, 548–555. https://doi.org/10.1109/BigData.2018.8622592

Psarra, S. (2005). Spatial Culture, Way-finding and Educational Message - The Impact of Layout on the Spatial, Social and Educational Experience of Visitors in Museums and Galleries. In S. Macleod (Ed.), *Reshaping museum space: architecture, design, exhibitions.* (pp. 78–94). Taylor and Francis.

Reda, K., Febretti, A., Knoll, A., Aurisano, J., Leigh, J., Johnson, A., Papka, M. E., & Hereld, M. (2013). Visualizing large, heterogeneous data in hybrid-reality environments. *IEEE Computer Graphics and Applications*, *33*(4), 38–48. https://doi.org/10.1109/MCG.2013.37

Ride with GPS | Bike Route Planner and Cycling Navigation App. (n.d.). Retrieved April 20, 2022, from https://ridewithgps.com/

Roberts, J. C., Ritsos, P. D., Badam, S. K., Brodbeck, D., Kennedy, J., & Elmqvist, N. (2014). Visualization beyond the desktop-the next big thing. *IEEE Computer Graphics and Applications*, *34*(6), 26–34. https://doi.org/10.1109/MCG.2014.82

Rogers, Y. (2006). Moving on from Weiser's Vision of Calm Computing: Engaging UbiComp Experiences. In *LNCS* (Vol. 4206).

Savage, N. (2010). Cycling through data. *Communications of the ACM*, *53*(9), 16. https://doi.org/10.1145/1810891.1810898

Savino, G. L., Meyer, L., Schade, E. E. S., Tenbrink, T., & Schöning, J. (2020, October 5). Point me in the right direction: Understanding user behaviour with as-the-crow-flies navigation. *Conference Proceedings - 22nd International Conference on Human-Computer Interaction with Mobile Devices and Services: Expanding the Horizon of Mobile Interaction, MobileHCI 2020*. https://doi.org/10.1145/3379503.3403539

Sayre, S., & Wetterlund, K. (2008). The Social Life of Technology for Museum Visitors. *Visual Arts Research*, *34*(2), 85–94. http://www.jstor.org/stable/20715478

Schemalsteig, D., Fuhrmann, A., Szavalari, Z., & Gervautz, M. (1996). Studierstube – An Environment for Collaboration in Augmented Reality. *In CVE '96 Workshop Proceedings*, 37–48.

Schilit, B., Adams, N., & Want, R. (1994). Context-Aware Computing Applications. *1994 First Workshop on Mobile Computing Systems and Applications*, 85–90. https://doi.org/10.1109/WMCSA.1994.16

Sears, A., Lin, M., Jacko, J., & Xiao, Y. (2003). When computers fade: Pervasive computing and situationally-induced impairments and disabilities. *HCI International*.

Shiro, K. (2019). *InterPoser : Visualizing Interpolated Movements for Bouldering Training*. 1–6.

Shneiderman, B. (1996). The eyes have it: a task by data type taxonomy for information visualizations. *Proceedings 1996 IEEE Symposium on Visual Languages*, 336–343. https://doi.org/10.1109/VL.1996.545307

SKIBA, P. F., CHIDNOK, W., VANHATALO, A., & JONES, A. M. (2012). Modeling the Expenditure and Reconstitution of Work Capacity above Critical Power. *Medicine & Science in Sports & Exercise*, *44*(8), 1526–1532. https://doi.org/10.1249/MSS.0b013e3182517a80

Skiba, P. F., & Clarke, D. C. (2021). The W′ Balance Model: Mathematical and Methodological Considerations. *International Journal of Sports Physiology and Performance*, *16*(11), 1561–1572. https://doi.org/10.1123/ijspp.2021-0205

Smaldone, S., Tonde, C., Ananthanarayanan, V. K., Elgammal, A., & Iftode, L. (2011). The cyber-physical bike. *Proceedings of the 12th Workshop on Mobile Computing Systems and Applications - HotMobile '11*, 56. https://doi.org/10.1145/2184489.2184502

Smith, M., Streeter, J., Burnett, G., & Gabbard, J. L. (2015). Visual Search Tasks: The Effects of Head-up Displays on Driving and Task Performance. *Proceedings of the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, 80–87. https://doi.org/10.1145/2799250.2799291

Snowdon, D., Greenhalgh, C., & Benford, S. (1995). What you see is not what I see: Subjectivity in virtual environments. *Framework for Immersive Virtual Environments (FIVE'95)*.

Star, S. L., & Griesemer, J. R. (1989). *Institutional Ecology, "Translations" and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39* (Vol. 19, Issue 3).

Stefik, M., Bobrow, D. G., Foster, G., Lanning, S., & Tatar, D. (1987). WYSIWIS revised: Early Experiences with Multiuser Interfaces. *ACM Transactions on Information Systems (TOIS)*, *5*(2), 147–167. https://doi.org/10.1145/27636.28056

Steltenpohl, H., & Bouwer, A. (2013). Vibrobelt. *Proceedings of the 2013 International Conference on Intelligent User Interfaces - IUI '13*, 417. https://doi.org/10.1145/2449396.2449450

Strava | Run and Cycling Tracking on the Social Network for Athletes. (n.d.). Retrieved April 20, 2022, from https://www.strava.com/

Strava Beacons. (n.d.). Retrieved May 25, 2022, from https://blog.strava.com/beacon/

Strava Challenges – Strava Support. (n.d.). Retrieved April 20, 2022, from https://support.strava.com/hc/en-us/articles/216919177-Strava-Challenges

Strava Metro Home. (n.d.). Retrieved May 1, 2022, from https://metro.strava.com/

Stusak, S., Tabard, A., Sauka, F., Khot, R. A., & Butz, A. (2014). Activity sculptures: Exploring the impact of physical visualizations on running activity. *IEEE Transactions on Visualization and Computer Graphics*, *20*(12), 2201–2210. https://doi.org/10.1109/TVCG.2014.2352953

Sullivan, G. M., Artino, A. R., & Jr. (2013). Analyzing and interpreting data from likert-type scales. *Journal of Graduate Medical Education*, *5*(4), 541–542. https://doi.org/10.4300/JGME-5-4-18

The Wireless Sensor Network Solution - THIS IS ANT. (n.d.). Retrieved June 13, 2022, from https://www.thisisant.com/

Tolmie, P., Benford, S., Greenhalgh, C., Rodden, T., & Reeves, S. (2014). Supporting group interactions in museum visiting. *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 1049–1059. https://doi.org/10.1145/2531602.2531619

Topliss, B. H., Pampel, S. M., Burnett, G., & Gabbard, J. L. (2019). Evaluating Head-Up Displays across Windshield Locations. *Proceedings of the 11th International Conference on Automotive User Interfaces and Interactive Vehicular Applications - AutomotiveUI '19*, 244–253. https://doi.org/10.1145/3342197.3344524

Track your Run - ASICS Runkeeper™ Running Tracker App. (n.d.). Retrieved May 1, 2022, from https://runkeeper.com/cms/

TrainingPeaks | Reclaim Your Race Day. (n.d.). Retrieved April 20, 2022, from https://www.trainingpeaks.com/

Trefzger, M., Blascheck, T., Raschke, M., Hausmann, S., & Schlegel, T. (2018). A visual comparison of gaze behavior from pedestrians and cyclists. *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications - ETRA '18*, 1–5. https://doi.org/10.1145/3204493.3214307

Twiss-Garrity Beth, & Fisher, M. (2007). Remixing Exhibits: Constructing participatory narratives with on-line tools to augment museum experiences. *Proceedings of Museums and the Web 2007*. http://www.archimuse.com/mw2007/papers/fisher/fisher.html

Vainstein, N., Kuflik, T., & Lanir, J. (2016). *Towards Using Mobile, Head-Worn Displays in Cultural Heritage*. https://doi.org/10.1145/2856767.2856802 van der Vaart, M., & Damala, A. (2015). Through the Loupe: Visitor engagement with a primarily text-based handheld AR application. *2015 Digital Heritage International Congress, Digital Heritage 2015*, 565–572. https://doi.org/10.1109/DigitalHeritage.2015.7419574

Vlahakis, V., Ioannidis, N., Karigiannis, J., Tsotros, M., Gounaris, M., Almeida, L., Stricker, D., Gleue, T., Christou, I. T., Carlucci, R., & Ioannidis, N. (2001). ARCHEOGUIDE: First results of an Augmented Reality, Mobile Computing System in Cultural Heritage Sites. *Proceedings of the 2001 Conference on Virtual Reality, Archeology, and Cultural Heritage - VAST '01*, *January*, 131. https://doi.org/10.1145/584993.585015

Vo, D. B., Saari, J., & Brewster, S. (2021, May 8). TactiHelm: Tactile Feedback in a Cycling Helmet for Collision Avoidance. *Conference on Human Factors in Computing Systems - Proceedings*. https://doi.org/10.1145/3411763.3451580

Waldron, M., Worsfold, P., White, C., & Murray, S. (2011). Swarming behaviour in elite race bunch cycling: A case study. *International Journal of Performance Analysis in Sport*, *11*(1), 14–25. https://doi.org/10.1080/24748668.2011.11868525

Walmink, W., Chatham, A., & Mueller, F. (2014). Interaction opportunities around helmet design. *Proceedings of the Extended Abstracts of the 32nd Annual ACM Conference on Human Factors in Computing Systems - CHI EA '14*, 367–370. https://doi.org/10.1145/2559206.2574803

Walmink, W., Wilde, D., & Mueller, F. "Floyd." (2014). *Displaying heart rate data on a bicycle helmet to support social exertion experiences*. 97–104. https://doi.org/10.1145/2540930.2540970

Warner, S., Dixon, M. A., & Chalip, L. (2012). THE IMPACT OF FORMAL VERSUS INFORMAL SPORT: MAPPING THE DIFFERENCES IN SENSE OF COMMUNITY. *Journal of Community Psychology*, *40*(8), 983–1003. https://doi.org/10.1002/jcop.21506

Weilenmann, A., Hillman, T., & Jungselius, B. (2013). Instagram at the Museum: Communicating the Museum Experience Through Social Photo Sharing. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1843–1852. https://doi.org/10.1145/2470654.2466243

Weiser, M., & Mark. (1993). Some computer science issues in ubiquitous computing. *Communications of the ACM*, *36*(7), 75–84. https://doi.org/10.1145/159544.159617

Winter, J. C. F. de, de Winter, J. C. F., & Dodou, D. (2010). Five-point Likert items: t test versus Mann-Whitney-Wilcoxon. *PRACTICAL ASSESSMENT, RESEARCH & EVALUATION*, 11. https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.650.3292

Wood, J. (2015). Visualizing Personal Progress in Participatory Sports Cycling Events. *IEEE Computer Graphics and Applications*, *35*(4), 73–81. https://doi.org/10.1109/MCG.2015.71

Woźniak, P. W., Dekker, L., Kiss, F., Velner, E., Kuijt, A., & Donker, S. F. (2020). Brotate and Tribike: Designing Smartphone Control for Cycling. *22nd International Conference on Human-Computer Interaction with Mobile Devices and Services*, 1–12. https://doi.org/10.1145/3379503.3405660

Wu, Y., Lan, J., Shu, X., Ji, C., Zhao, K., Wang, J., & Zhang, H. (2018). ITTVis: Interactive Visualization of Table Tennis Data. *IEEE Transactions on Visualization and Computer Graphics*, *24*(1), 709–718. https://doi.org/10.1109/TVCG.2017.2744218

Zhu, H. (2004). From WYSIWIS to WISINWIS: Role-based collaboration. *Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics*, *6*, 5441–5446. https://doi.org/10.1109/ICSMC.2004.1401059

Zhu, H. (2006). Role mechanisms in collaborative systems. *International Journal of Production Research*, *44*(1), 181–193. https://doi.org/10.1080/00207540500247495

Zwahlen, H. T., Adams, C. C., & DeBals, D. P. (1988). Safety Aspects of CRT Touch Panel Controls in Automobiles. *Vision in Vehicles II. Proceedings of the Second International Conference on Vision in Vehicles*, 335–344.

APPENDIX A – MUSEUM STUDY'S BRIEFING SCRIPT

Part 1: Introduction to the exhibit

Hello and welcome to the Halifax Explosion Exhibit. My name is [investigator], let me take few minutes to explain some basic interaction methods and safety features before you begin.

To fully interact with this exhibit one of you will need to use a HoloLens, which is a head-mounted display capable of showing holograms in front of you [shows HoloLens]. With it you will be able to see 3D models of key buildings of Halifax rendered above the wood representation of the city.

Interacting with the virtual content is fairly simple, first the only input method is to press this button, [gives button to participant] with it you can enlarge the virtual buildings. To do so, you look at a specific building that you want to explore in more detail, placing the cursor on top of it, and press the button. The building will then enlarge in the center of the table, half a meter above its surface. You can make the building small again by looking at it and clicking again or by enlarging another building.

Some buildings have images inside of that that can be seen moving towards them (or into them) while they are enlarged

The HoloLens can track itself around the table, so it always knows where it is regarding the wood model. Therefore, we encourage you to move around freely and explore the exhibit from different angles.

Be aware that you may become really focused on the virtual content presented over the real world and may become less aware of your surroundings (something we refer as immersion). Please try to remember that there might other people around you while you are moving around the wood model.

When you are finished exploring the virtual elements of the exhibit please let me know so I can take the HoloLens safely from you.

Lastly, AR headsets generally do not cause nausea or disorientation however if you feel discomfort of any kind please let us know so we can assist you to remove the headset.

Finally, have fun! Do you have any questions before we proceed? [answer any questions]

Part 2: Invitation to participate in the study

We are also running an exploratory study looking at how people interact with this mixed setup of virtual and real content and on how they communicate what they currently seeing with other. You are free to explore the exhibit whether or not you wish to participate in the study, but if you do agree to participate we will need your consent to record video and audio of both of you while you explore this exhibit and to ask a couple questions in the end. Your faces will be blurred in the video to preserve anonymity.

After the study we will provide you with a debriefing form that has all the information about the study including who to contact if you have questions and how to withdraw your data from the study if you change your mind about participating. Do we have your consent to collect data during the study?

APPENDIX B – SYSTEM USABILITY SCORE (SUS)

I think that I would like to use this system frequently	1 Strongly Disagree	2	3	4	5 Strongly Agree
I found the system unnecessarily complex	1 Strongly Disagree	2	3	4	5 Strongly Agree
I thought the system was easy to use	1 Strongly Disagree	2	3	4	5 Strongly Agree
I think that I would need the support of a technical person to be able to use this system	1 Strongly Disagree	2	3	4	5 Strongly Agree
I found the various functions in this system were well integrated	1 Strongly Disagree	2	3	4	5 Strongly Agree
I thought there was too much inconsistency in this system	1 Strongly Disagree	2	3	4	5 Strongly Agree

I would imagine that most people would learn to use this system very quickly	1 Strongly Disagree	2	3	4	5 Strongly Agree
I found the system very cumbersome to use	1 Strongly Disagree	2	3	4	5 Strongly Agree
I felt very confident using the system	1 Strongly Disagree	2	3	4	5 Strongly Agree
I needed to learn a lot of things before I could get going with this system	1 Strongly Disagree	2	3	4	5 Strongly Agree

APPENDIX C – MUSEUM STUDY'S SEMI-STRUCTURED INTERVIEW QUESTIONS

1. Do you think this technology can be used in other museum exhibits?
2. Were the two of you [three, four, …] able to communicate while exploring this exhibit? Were you [look at no HoloLens participant] able to understand what was presented in the virtual world?
3. What did you talked about while referring to the virtual content?
4. Do you have any other thoughts or comments about your experience?

APPENDIX D1 – MUSEUM STUDY'S POST-SESSION QUESTIONNAIRE (HOLOLENS USER)

It was difficult for me to communicate with others while using the HoloLens.	1 Strongly Disagree	2	3	4	5 Strongly Agree
It was easy for me to share and explain what I was seeing.	1 Strongly Disagree	2	3	4	5 Strongly Agree
I feel that my friends/family were disconnected from me while I was exploring the virtual content.	1 Strongly Disagree	2	3	4	5 Strongly Agree
I enjoyed this virtual experience combined with the physical model.	1 Strongly Disagree	2	3	4	5 Strongly Agree

APPENDIX D2 – MUSEUM STUDY'S POST-SESSION QUESTIONNAIRE (NO HOLOLENS USER)

It was difficult for me to communicate with the person using the HoloLens	1 Strongly Disagree	2	3	4	5 Strongly Agree
It was easy for me to understand what was being displayed in the virtual environment.	1 Strongly Disagree	2	3	4	5 Strongly Agree
I feel that I was disconnected from the person exploring the virtual content with the HoloLens	1 Strongly Disagree	2	3	4	5 Strongly Agree
I enjoyed this shared mixed-reality experience even without the headset.	1 Strongly Disagree	2	3	4	5 Strongly Agree

APPENDIX E – CONTEXTUAL INQUIRY CULTURAL MODEL

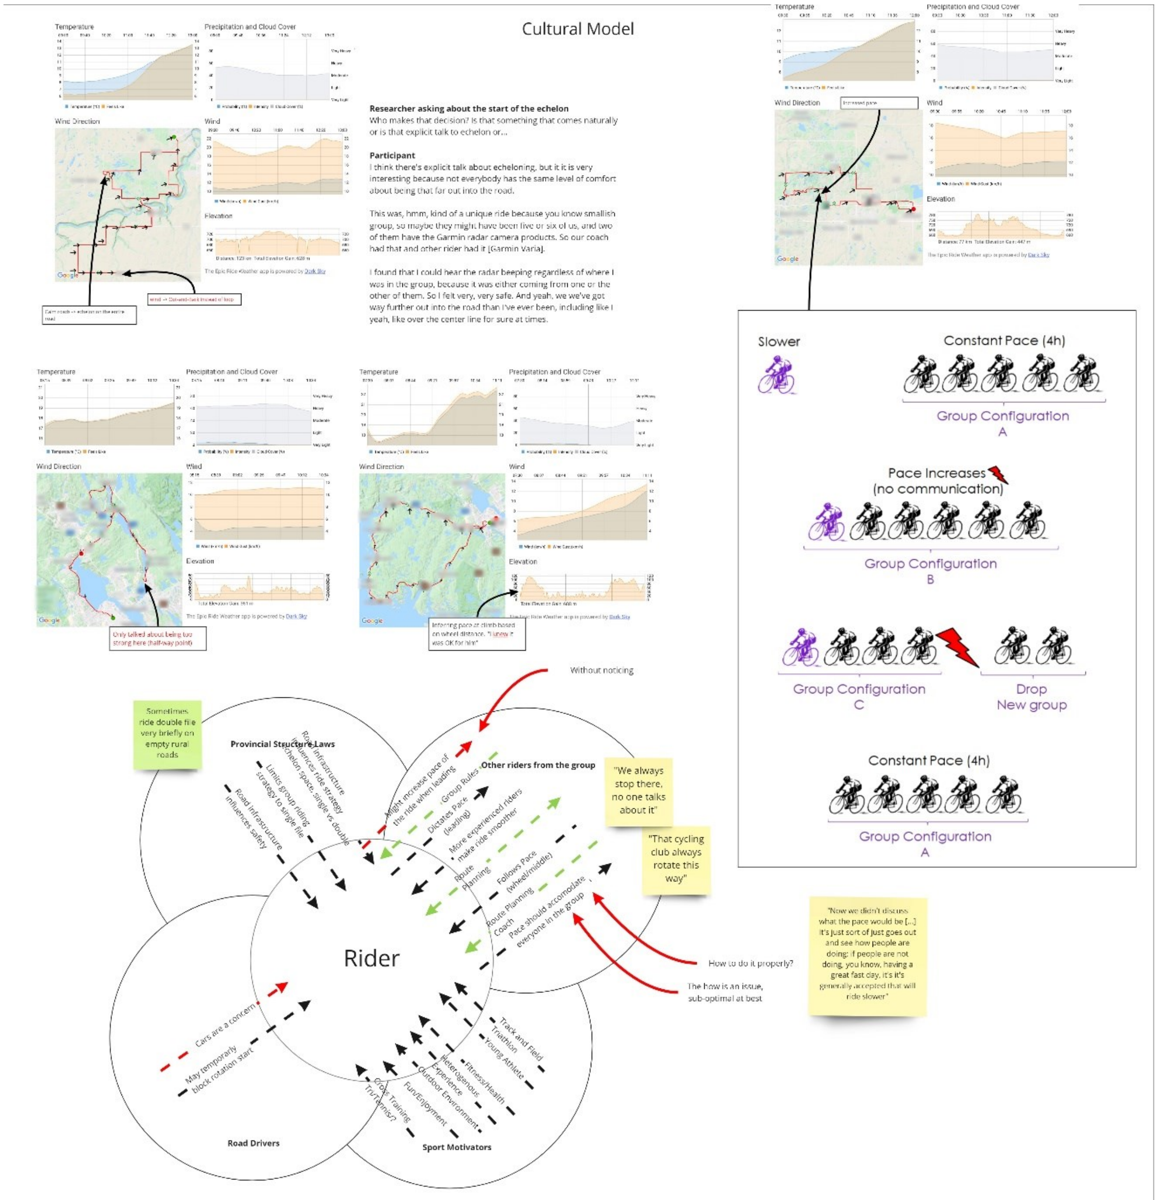

APPENDIX F – CONTEXTUAL INQUIRY FLOW MODEL

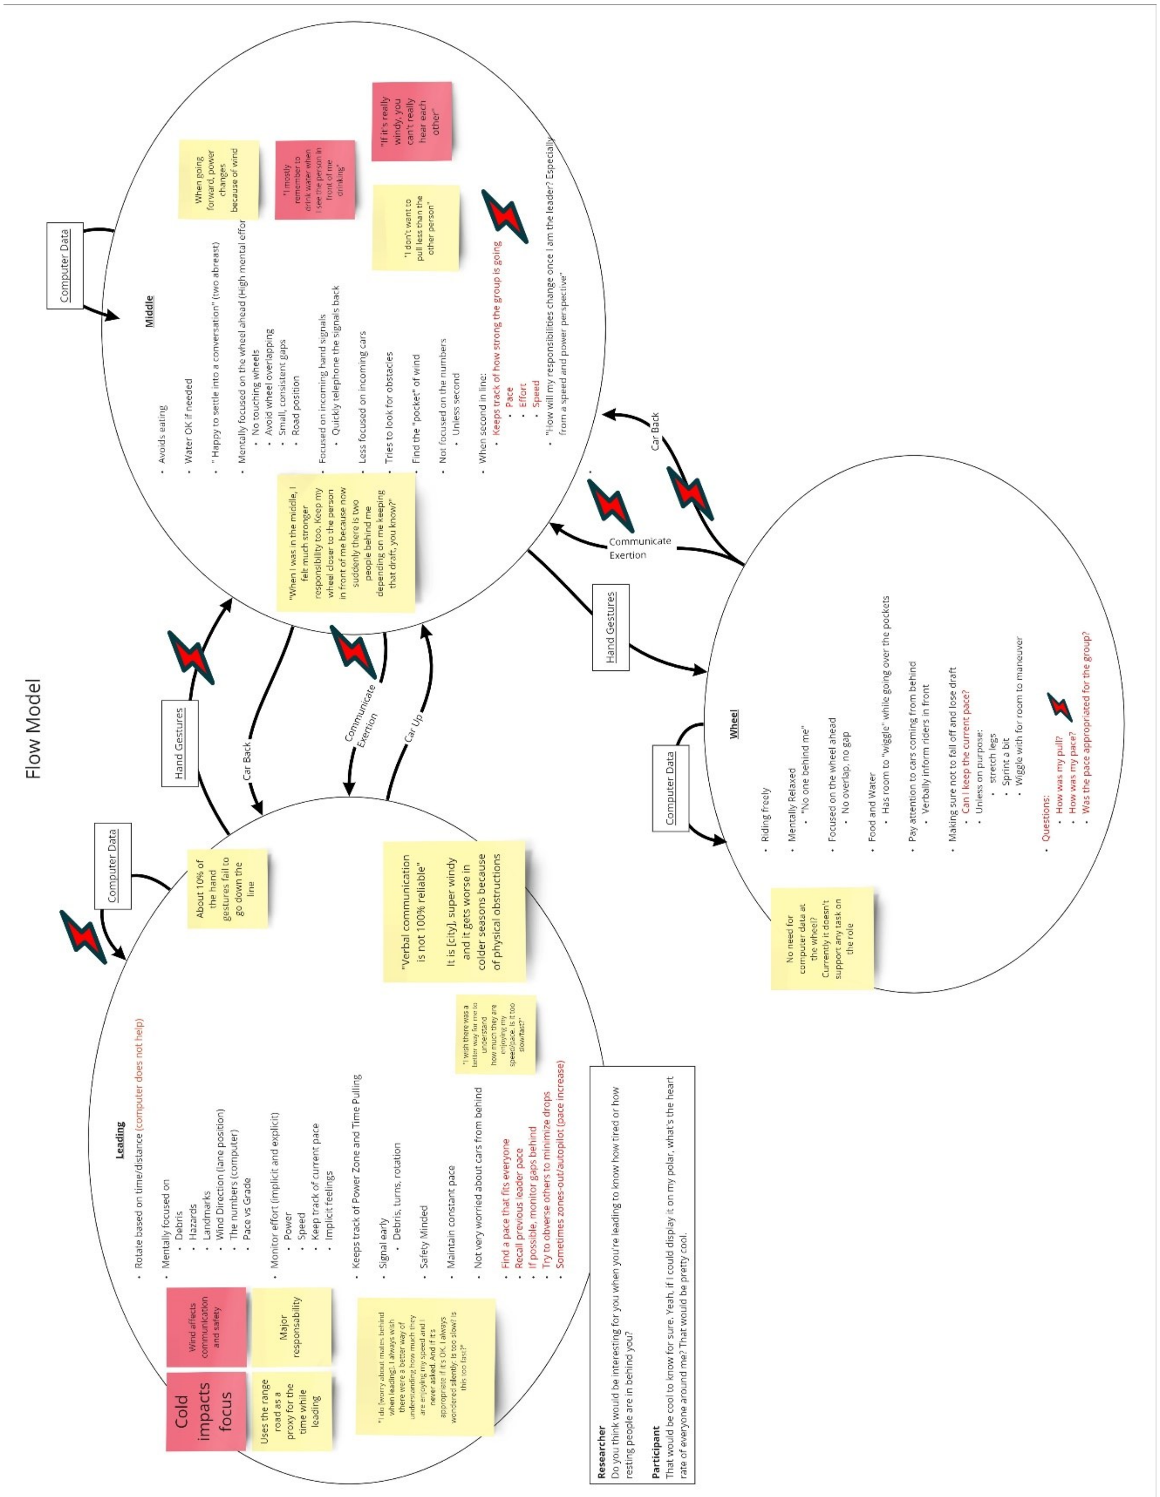

APPENDIX G – CONTEXTUAL INQUIRY ARTIFACT MODEL

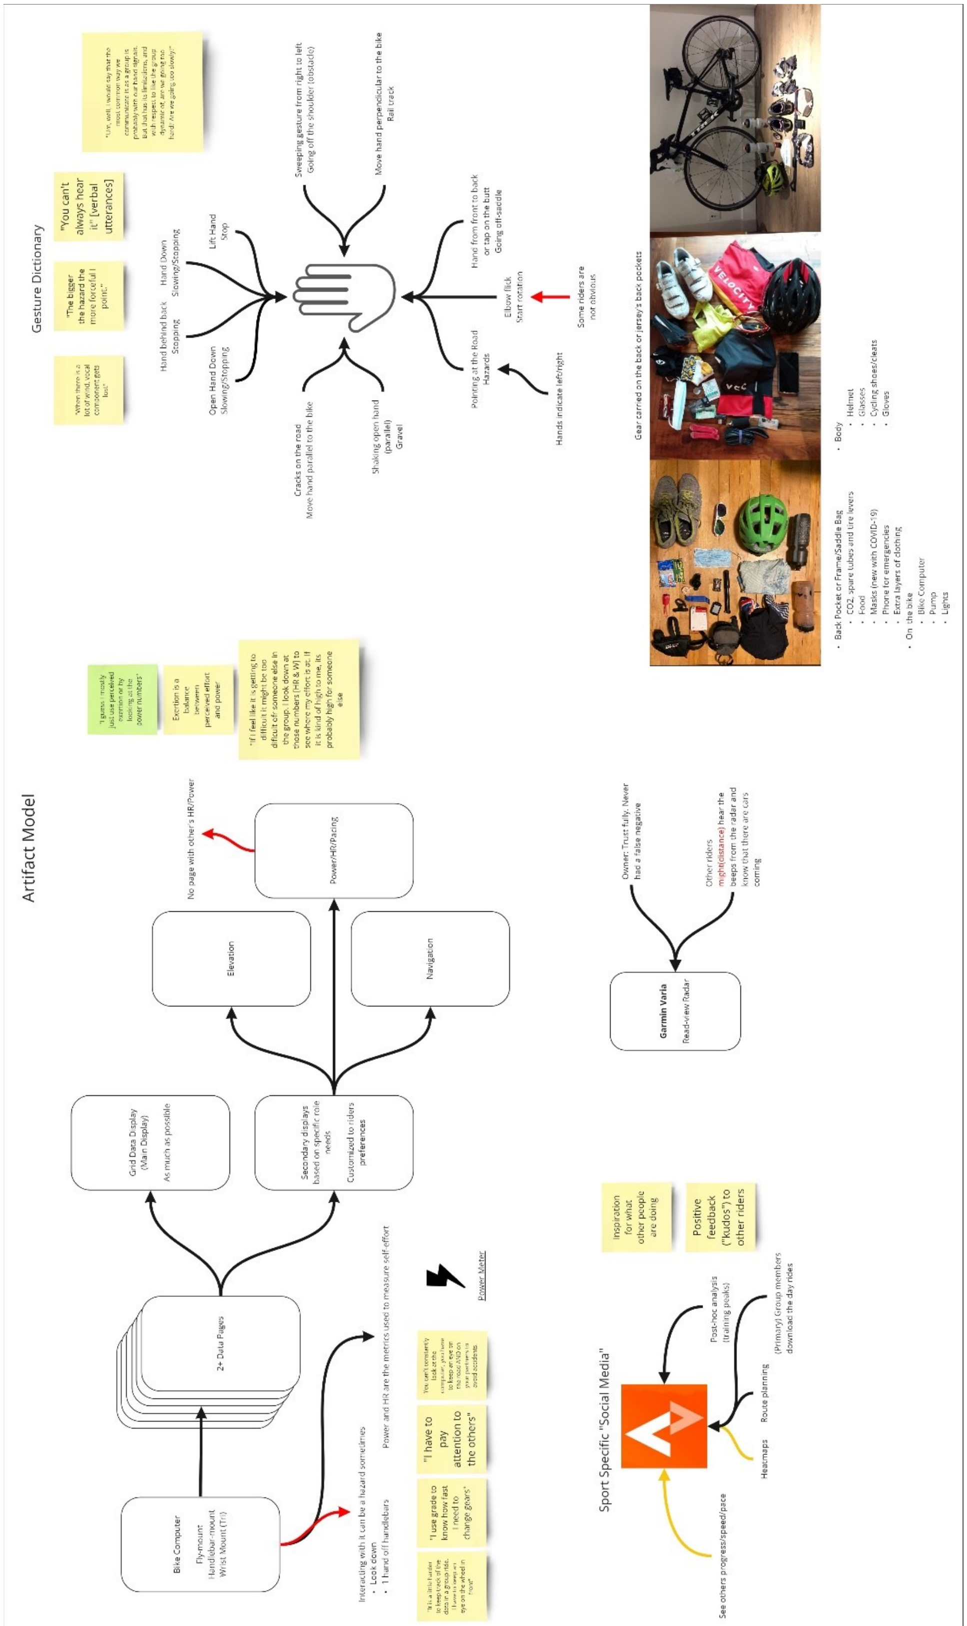

APPENDIX H – CONTEXTUAL INQUIRY PHYSICAL MODEL

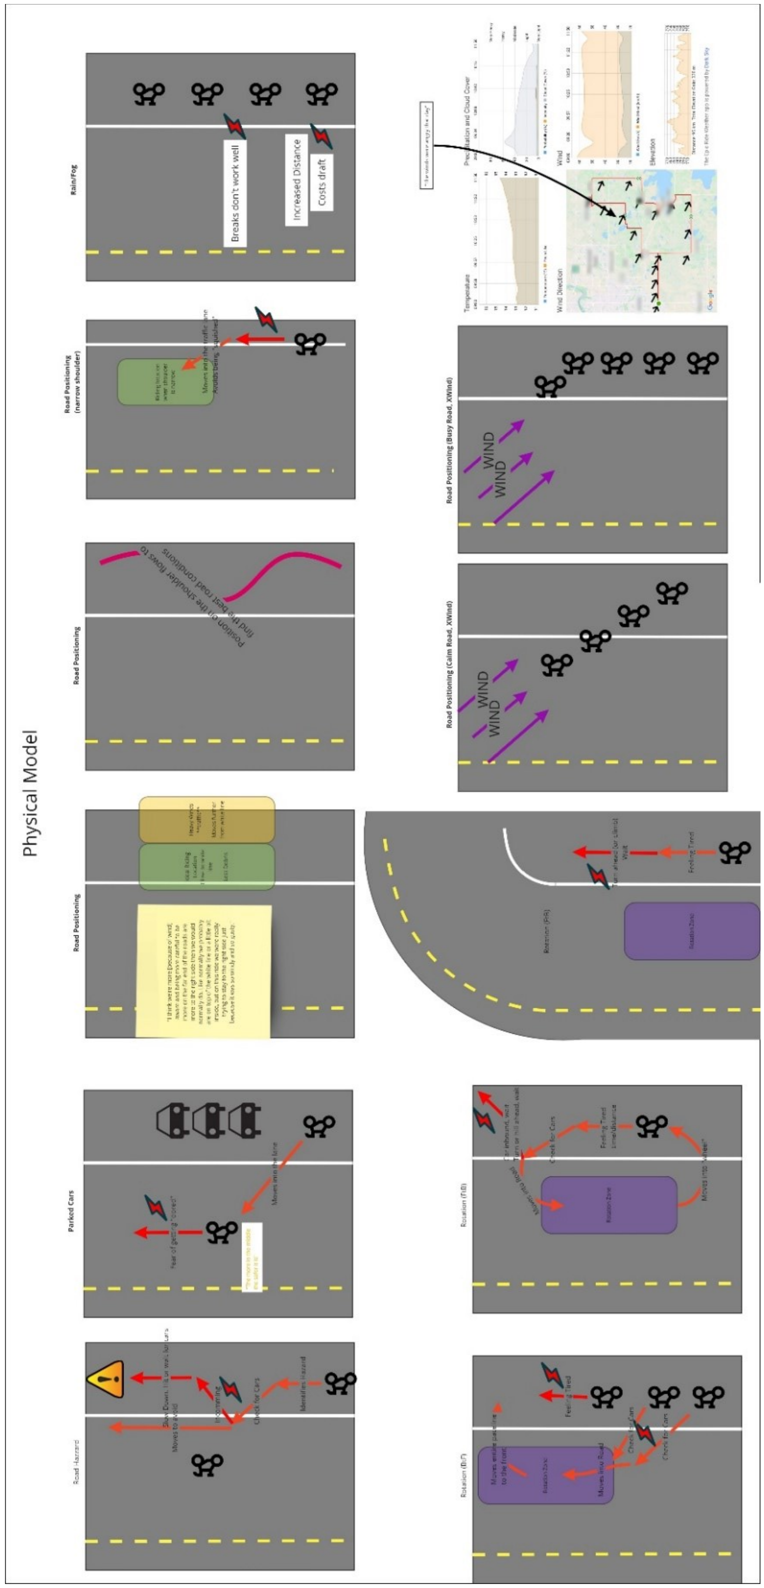

APPENDIX I1 – CONTEXTUAL INQUIRY SEQUENCE MODEL

APPENDIX I2 – CONTEXTUAL INQUIRY SEQUENCE MODEL (CONT)

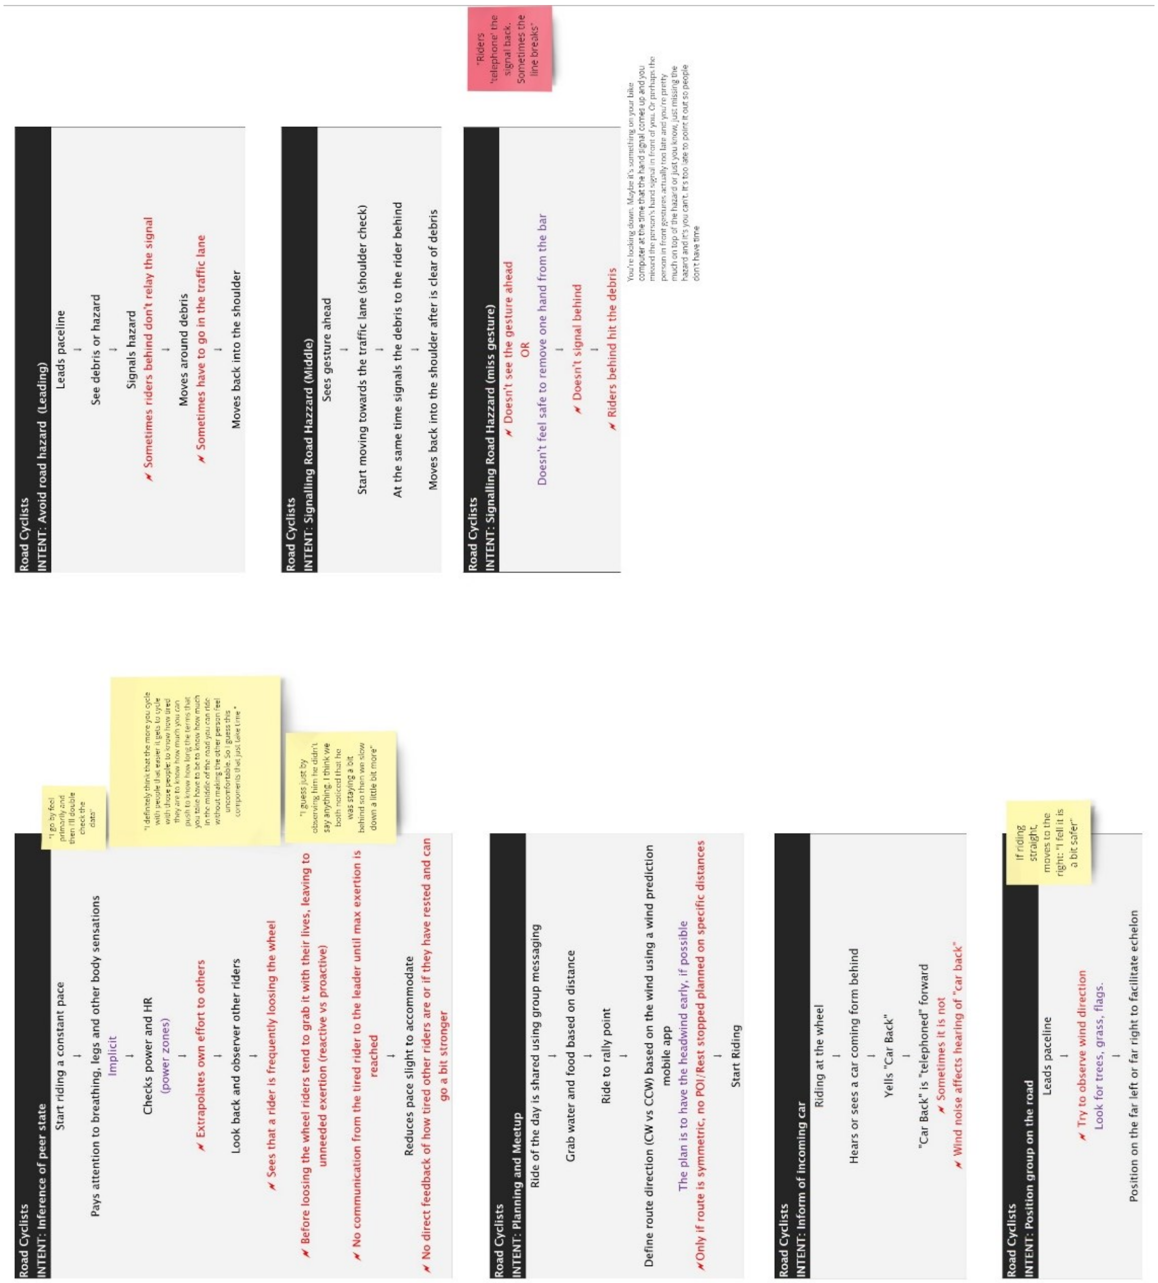

APPENDIX J – CONTEXTUAL INQUIRY STUDY ETHICS APPROVAL

DALHOUSIE UNIVERSITY
Research Services

Health Sciences Research Ethics Board
Letter of Approval
June 24, 2020
Juliano Franz
Computer Science\Computer Science

Dear Juliano,

REB #: 2020-5178
Project Title: Modeling Distributed Cognition in Cycling Packs

Effective Date: June 24, 2020
Expiry Date: June 24, 2021

The Health Sciences Research Ethics Board has reviewed your application for research involving humans and found the proposed research to be in accordance with the Tri-Council Policy Statement on *Ethical Conduct for Research Involving Humans.* This approval will be in effect for 12 months as indicated above. This approval is subject to the conditions listed below which constitute your on-going responsibilities with respect to the ethical conduct of this research.

Effective March 16, 2020: Notwithstanding this approval, any research conducted during the COVID-19 public health emergency must comply with federal and provincial public health advice as well as directives issued by Dalhousie University (or other facilities where the research will occur) regarding preventing the spread of COVID-19.

Sincerely,

[redacted]

Dr. Lori Weeks, Chair

APPENDIX K – ZWIFT STUDY ETHICS APPROVAL #1 (SOCIAL SCIENCES)

DALHOUSIE UNIVERSITY
Research Services

Social Sciences & Humanities Research Ethics Board
Letter of Approval

September 01, 2021
Juliano Franz
Computer Science\Computer Science

Dear Juliano,

REB #: 2021-5711
Project Title: Impacts of Sharing Performance Metrics on Group Road Cyclists

Effective Date: September 01, 2021
Expiry Date: September 01, 2022

The Social Sciences & Humanities Research Ethics Board has reviewed your application for research involving humans and found the proposed research to be in accordance with the Tri-Council Policy Statement on *Ethical Conduct for Research Involving Humans.* This approval will be in effect for 12 months as indicated above. This approval is subject to the conditions listed below which constitute your on-going responsibilities with respect to the ethical conduct of this research.

Effective March 16, 2020: Notwithstanding this approval, any research conducted during the COVID-19 public health emergency must comply with federal and provincial public health advice as well as directives from Dalhousie University (and/or other facilities or jurisdictions where the research will occur) regarding preventing the spread of COVID-19.

Sincerely,

[redacted]

Dr. Karen Foster, Chair

APPENDIX L – ZWIFT STUDY ETHICS APPROVAL #2 (HEALTH SCIENCES)

DALHOUSIE UNIVERSITY
Research Services

Health Sciences Research Ethics Board
Letter of Approval

September 01, 2021

Juliano Franz
Computer Science\Computer Science

Dear Juliano,

REB #: 2021-5711
Project Title: Impacts of Sharing Performance Metrics on Group Road Cyclists

Effective Date: September 01, 2021
Expiry Date: September 01, 2022

The Health Sciences Research Ethics Board has reviewed your application for research involving humans and found the proposed research to be in accordance with the Tri-Council Policy Statement on *Ethical Conduct for Research Involving Humans.* This approval will be in effect for 12 months as indicated above. This approval is subject to the conditions listed below which constitute your on-going responsibilities with respect to the ethical conduct of this research.

Effective March 16, 2020: Notwithstanding this approval, any research conducted during the COVID-19 public health emergency must comply with federal and provincial public health advice as well as directives from Dalhousie University (and/or other facilities or jurisdictions where the research will occur) regarding preventing the spread of COVID-19.

Sincerely,

[redacted]

Dr. Lori Weeks, Chair

APPENDIX M – BORG'S RPE TABLE

Table 12 - Table from (BORG, 1982)

How you might describe your exertion	**Borg rating of your exertion**	**Examples**
None	6	Reading a book, watching TV
Very, very light	7 to 8	Tying shoes
Very light	9 t0 10	Chores, like folding clothes, that seem to take little effort
Fairly light	11 to 12	Walking through the grocery store or other activities that require some effort but not enough to speed up your breathing
Somewhat hard	13 to 14	Brisk walking or other activities that require moderate effort and speed your heart rate a breathing but don't make you out of breath
Hard	15 to 16	Bicycling, swimming, or other activities that take vigorous effort and get the heart pounding and make breathing very fast
Very hard	17 to 18	The highest level of activity you can sustain
Very, very hard	19 to 20	A finishing kick in a race or other burst of activity that you can't maintain for long

APPENDIX N – MUSEUM STUDY ETHICS APPROVAL

DALHOUSIE UNIVERSITY
Research Services

Social Sciences & Humanities Research Ethics Board
Letter of Approval

June 22, 2018

Juliano Franz
Computer Science\Computer Science

Dear Juliano,

REB #: 2018-4519
Project Title: Sharing Mixed Reality Content in Museum Exhibits

Effective Date: June 22, 2018
Expiry Date: June 22, 2019

The Social Sciences & Humanities Research Ethics Board has reviewed your application for research involving humans and found the proposed research to be in accordance with the Tri-Council Policy Statement on *Ethical Conduct for Research Involving Humans.* This approval will be in effect for 12 months as indicated above. This approval is subject to the conditions listed below which constitute your on-going responsibilities with respect to the ethical conduct of this research.

Sincerely,

[redacted]

Dr. Karen Beazley, Chair

APPENDIX O – ANATOMY OF A BICYCLE

APPENDIX P – ZWIFT STUDY RIDE PARTICIPATION PROTOCOL

Prior to each ride, participants would answer a doodle pool to choose the time for their session at a time that all three members could participate. There were no restrictions on time, except between midnight and 6 am on the researcher time-zone (UTC-4). Once a time was agreed upon, participants received a Zwift Meetup[30] invite (through Zwift's application) for the day of the ride containing start time and the route (Ocean Cliffside Loop).

At the day of the trial, participants would join Discord so they could communicate with the researcher and each other while riding. At the same time, join the meetup on Zwift (already on their bicycles wearing riding clothes). Once all riders connected to the meetup, the following steps were followed:

1. The researcher would remind everyone of the ride instructions.
2. Participants would keep spinning waiting for the start time.
3. Ride starts.
4. Participants start riding in a paceline.
5. At regular intervals the researcher would ask the participant's RPE values
    a. Researcher "deafens" all other participants.
    b. Ask the exertion question.
    c. Updates RPE values
    d. Repeat process to other riders.

[30] https://support.zwift.com/en_us/meetups-HJP7iUd4r

    e. Once all riders updated their RPE values the researcher "undefens" all

6. At every five minutes, researcher asks group to rotate.
7. Ride finishes.
8. Participants receive survey link via e-mail so they can answer before closing for the day.

APPENDIX O – MUSEUM STUDY DEBRIEFING SCRIPT

Thank you for participating in our study! We are trying to understand how people interact with the augmented reality world and how they share their experience with others. Participation in our study was voluntary. If you would like to withdraw your data, for any reason, you can do so by contacting Juliano Franz (juliano.franz@dal.ca) and include your participant ID listed above. You do not have to provide a reason for withdrawing your data.

We asked you to explore the mixed reality content (virtual elements inserted into the real world) freely, so we could observe how you and other participants interact with the setting. While you explored the virtual content and shared your experience with other participants, we recorded video and audio, and logged the position and orientation of the headset. All data will be kept confidential. We use numerical ids to preserve anonymity of all the data we record, and your face will be blurred in the videos.

We are happy to talk to you about any questions or concerns you may have about your participation in this research study. Please contact Juliano Franz (juliano.franz@dal.ca) at any time with questions, comments, or concerns about the research study. If you have any ethical concerns about your participation in this research, you may also contact Research Ethics, Dalhousie University at (902) 494-1462, or email: ethics@dal.ca (and reference REB file # 2018-4519)