



Williamson, J. R., Li, J., Vinayagamoorthy, V., Shamma, D. A. and Cesar, P. (2021) Proxemics and Social Interactions in an Instrumented Virtual Reality Workshop. In: ACM Conferences on Human Computer Interaction, Yokohama, Japan, 08-12 May 2021, p. 253. ISBN 9781450380966.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

© Association for Computing Machinery 2021. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in ACM Conferences on Human Computer Interaction, Yokohama, Japan, 08-12 May 2021, p. 253. ISBN 9781450380966. <http://dx.doi.org/10.1145/3411764.3445729>.

<http://eprints.gla.ac.uk/234460/>

Deposited on: 11 June 2021

Proxemics and Social Interactions in an Instrumented Virtual Reality Workshop

JULIE R. WILLIAMSON, University of Glasgow, Scotland

JIE LI, CWI, Netherlands

VINOBA VINAYAGAMOORTHY, BBC R&D, United Kingdom

DAVID A. SHAMMA, CWI, Netherlands

PABLO CESAR, CWI, Netherlands




Fig. 1. An overhead view of the main room showing participants locations and viewports while they attend the workshop.

Virtual environments (VEs) can create collaborative and social spaces, which are increasingly important in the face of remote work and travel reduction. Recent advances, such as more open and widely available platforms, create new possibilities to observe and analyse interaction in VEs. Using a custom instrumented build of Mozilla Hubs to measure position and orientation, we conducted an academic workshop to facilitate a range of typical workshop activities. We analysed social interactions during a keynote, small group breakouts, and informal networking/hallway conversations. Our mixed-methods approach combined environment logging, observations, and semi-structured interviews. The results demonstrate how small and large spaces influenced group formation, shared attention, and personal space, where smaller rooms facilitated more cohesive groups while larger rooms made small group formation challenging but personal space more flexible. Beyond our findings, we show how the combination of data and insights can fuel collaborative spaces' design and deliver more effective virtual workshops.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

Manuscript submitted to ACM

CCS Concepts: • **Human-centered computing** → **Virtual reality**; **Collaborative interaction**; **Field studies**.

Additional Key Words and Phrases: Virtual Environments, Virtual Meetings, Social Signal Processing, Interviews.

ACM Reference Format:

Julie R. Williamson, Jie Li, Vinoba Vinayagamoorthy, David A. Shamma, and Pablo Cesar. 2020. Proxemics and Social Interactions in an Instrumented Virtual Reality Workshop. In *CHI '21: ACM Computer Human Interaction, May 8–13, 2021, Yokohama, Japan*. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Virtual Environments (VEs) are not a replacement for face-to-face interaction, but VEs can provide solutions in the face of increasing pressures to reduce travel and work remotely. In this paper, we explore how virtual academic workshops can be effective at enabling social group activities like keynote talks, small group discussions, and informal networking in a VE. Taking advantage of customisation in the open source VE Mozilla Hubs¹, we built a custom environment instrumented with position and movement logging. We designed and evaluated how the workshop programme, the virtual spaces, and hosting approaches contributed to a successful event.

Current VEs vary widely in affordances, fidelity, scale, and accessibility. VEs like Mozilla Hubs, Gather Town², and VRChat³ all create social VEs, but result in dramatically different experiences. Mozilla Hubs and VRChat provide fully 3D environments that can be experienced on a desktop or using a head-mounted display (HMD). Gather Town uses a 2D map, but incorporates video conferencing for groups to chat. VRChat is a massively online multiplayer environment, averaging over 10,000 users daily. In contrast, Mozilla Hubs and Gather Town support a maximum of 25 and 50 users respectively, although premium Gather Town rooms can host up to 500. VRChat requires a relatively high specification PC suitable for gaming, but Mozilla Hubs and Gather Town run in most standard browsers. Accessibility in VEs presents an on-going challenge [50], and open source VEs like Mozilla Hubs create the best opportunities to address accessibility requirements. This is especially important for scholarly community events which must prioritise accessibility and inclusion. Organisers designing virtual events must balance complex issues and often competing needs [8], making it challenging to choose the right tools for the experience they are trying to create. For an academic workshop, our priorities were to facilitate social interaction in small groups, create opportunities to network, and facilitate discussion on open-ended topics. As researchers in social virtual reality (VR), we wanted to organise an event to discuss the open issues of social VR in social VR. Although we initially intended this as a hybrid event with physical and virtual spaces, the pandemic of 2020 meant we had to rapidly adapt to a fully virtual event.

In this paper, we present research on running an academic workshop in Mozilla Hubs. The workshop programme included a keynote talk, individual pitches, small group breakouts, and information breaks/networking opportunities. We designed the virtual environments to support these activities using a large outdoor environment for collective activities (see Figure 1) and breakout rooms for smaller groups. We customised a Mozilla Hubs Cloud instance to log user positions and orientations, generating a proxemic dataset from the workshop event. Combining a proxemic dataset with observations and post-workshop interviews, we explore whether proxemic interactions in a VE can be analysed using metaphors from physical settings, what is the relationship between group activity and VEs, and what social cues are available or needed? To answer these questions, we used an open source VE instrumented to support quantitative analysis of position, orientation, framerate, and user input as social signals, and release the code

¹Mozilla Hubs is an open source platform for hosting virtual events.

²Gather Town combines 2D maps with video conferencing.

³VRChat is an online massively multiplayer social environment.

for collection, analysis, and captured dataset (see Appendix A). We demonstrate the use of quantitative analysis for measuring proxemics, interactions, and personal space, inspired by interaction in physical settings. We combine this with a mixed methods analysis of group formation and social cues to see how the environment design affects small and large groups collaborating in a VE. Finally, we conclude with a discussion of open challenges for social group interaction in academic workshops in VEs, including dynamics like expressing oneself effectively, knowing where to stand, and interrupting existing conversations.

2 BACKGROUND

There is a substantial body of work on socialising and collaborating in VEs that informed our approach. We discuss the existing methods for evaluations in VEs, and how our instrumented approach adds insights to these methods. We also draw from literature on proxemics and personal space, discussing how these concepts from physical environments translate to the virtual. Finally, we discuss research on group formation and social regulation in VEs and their use on informing the design of social activities.

2.1 Methods for Analysing Behaviour in Virtual Environments

Previous findings on proxemics, spatial positioning, and group interaction in VEs are primarily supported by observational data and questionnaires. Metrics based on standardised questionnaires inform us about the perceived presence and immersion [45, 55]. More recently, an extension to these questionnaires aims at measuring the quality of the interaction of participants in a social VR experience [28]. These methods require post self-reflection from the users and are not suitable for capturing actual movement and behaviour of users. Ethnographic methods based on observations have been used to better understand user behaviour. For example, observational notes can be used to analyse behaviour [27]. Additional examples include the study of user behaviour in immersive games like 'there' [9] or World of Warcraft [14], and in large virtual environments like Second Life [56]. Audio and voice have also been captured for analysing interaction in VEs [7, 12, 51].

There are limited examples of more systematic instrumentation of VEs for analysing user position and movement. For example, Friedman et al. [17] explores the spatial social behaviour in Second Life based on software bots that wander around and collect data, while Ranathunga et al. [38] reports an infrastructure for identifying events and happenings in Second Life. More recently, Le et al. [27] perform quantitative analysis of the behaviour of participants in a virtual conference based on the logs from a voice and chat application (Discord) integrated in a virtual environment (Mozilla Hubs).

VEs are growing in popularity, influencing training, collaboration, social experiences [44]. Previous work on scientific events has explored large-scale gatherings [16] and the affordances of the existing technology [34, 42], in some cases with limited instrumentation of the space or the auxiliary communication channels [27]. In contrast, our work is focused on understanding the individual and group behaviour of participants attending a small and focused scientific event in a VE. Our method is inspired by interaction in physical settings, for example using cameras [11], mobile devices [13, 18], and wearable sensors [26] to capture high resolution data on position and orientation. We also completed observations and post-workshop interviews to gather qualitative insights. Thus, our approach uses mixed-methods, drawing from both qualitative and quantitative techniques to analyse social interaction in a VE.

2.2 Proximity and Personal Space

The flow and changes in interpersonal distances between individuals in a shared space is an integral part of nonverbal communication. Neither the intention nor awareness of specific behavioural cues is necessary and yet these cues dictate part of the conversation. In fact, it adds to the richness and versatility of the conversation [2]. Hall's theoretical model of proxemics, largely based on middle class American adults, focused on four concentric zones to define different levels of social intimacy which 'informed' the interpersonal distance maintained in between people [19]. These 'circles' were *intimate* (<0.45m), *personal* (0.45m-1.2m), *social* (1.2m-3.6m) and *public* (>3.6m). However, there are many factors which affect the proximal relationship between people including age, culture, environmental context, interpersonal relationship and emotional state of the individuals [2, 4, 20, 40, 48]. Proximal behaviour is also interconnected with other modalities of expression such as gaze behaviour. Argyle and Dean's [3] theory of equilibrium posits that appropriate interpersonal distance is maintained between individuals by balancing increasing proximity with reduced eye contact. The lower limits of the distance is determined by physical contact whereas the upper limit is defined by factor of visibility and audibility [19].

VEs are often crafted as representations of physical spaces. However, it is unclear how much of our understanding of physical spaces transfers to analysing virtual spaces. Users may enter a shared virtual space with physical metaphors in mind, but in the absence of physical cues and constraints do these metaphors transfer across? Hecht et al. [21] reported that the shape of personal space in the real world was remarkably close to a circular zone with a radius of about one meter and this was similar to the personal distance participants maintained to an avatar in a virtual space. Bailenson et al [5] tested the equilibrium theory in an immersive VE using two objective measures—minimum distance participants maintained while approaching a virtual male agent and invasion duration within the agents' intimate space. They reported that participants avoided direct contact with the agent and respected the personal space of the agent (relative to the control condition: a cylinder). There were also gender differences with female participants responding to mutual gaze behaviour with the agent more than male participants. Bailenson et al. [6] extended their work to include head gestures, female virtual agents and perceived agency. Again, participants maintained greater personal distance with the agent that engaged in mutual gaze in addition to getting out of the way when the virtual agent approached the participants. This effect was also noted by Llobera et al. [29] where they noted that the shorter the distance between virtual agents and the participant, the greater the physiological arousal as indicated through electrodermal activity. Wilcox et al. [54] further reported emotional discomfort in situations of close encounters, both through self-reported means and increased skin conductance. Yee et al. [56] conducted a observational study in Second Life to explore if social norms of gender, interpersonal distance and gaze behaviour transferred to VEs even though participants used a keyboard and mouse as opposed to their bodies. They found support for the equilibrium theory through avoidance of collision—a result reported by Friedman et al. [17] as well. Additionally, results showed that male dyads maintained larger interpersonal distances and less mutual gaze than female dyads. They also reported that male dyads were significantly less likely to maintain mutual gaze in indoor locations which made sense due to the implication of higher levels of intimacy associated with more mutual gaze and the constraints placed on maintaining larger interpersonal distance in smaller spaces [56]. In a similar vein, Bonsch et al. [10] reported that participants chose to maintain larger interpersonal distances to virtual agents perceived to be angry in comparison to happy virtual agents.

"Personal space" is clearly important in both physical and virtual worlds. However, it's also clear that we don't have a complete picture of proxemics and personal space when applied to virtual spaces. For instance, unlike Bailenson et al. [5], Llobera et al. [29] didn't detect differences in participant physiological responses when they were approached by

cylinders. Takahashi et al. [49] proposed that cones facing an observer were perceived as socially discomforting and threatening. Nearly two decades ago, Bailenson et al. [5, 6, 56] suggested that proxemics could prove a valuable gauge for measuring the behavioural realism of virtual agents/avatars. Recently, McVeigh-Schultz et al. [30] laid out some design considerations, extracted from expert interviews with creators of VR applications, which might shape social interactions in shared virtual spaces including awareness of affordances of proxemics and personal space management in order to tackle issues of harassment. Positional data can now be easily surfaced in VR platforms allowing us to use proxemics as part of a mixed-method evaluative toolkit to, at the very least study, user experience in a bid to design better spaces.

2.3 Forming Groups and Social Activities

As with physical environments, the design of VEs influences how people form groups and join activities. Inspired by urban ethnography, Moore et al. describe environment design in terms of accessibility, social density, activity resources, and hosts [31]. For example, large open spaces can be difficult to fill with enough people to achieve “social density.” Virtual spaces that are not restricted by physical building or space constraints often fail to create cosy places where social contact comes easily.

Apart from influences of environmental designs, social cohesion and social identity have long been recognised as two main factors that impact on group formation [23]. Social cohesion traces group formation to processes of interpersonal attraction [24], while the social identity defines by considering identification, or self-categorisation to be the mechanism of group formation [23]. Postmes et al. [36] found that, in computer-mediated communication, where people do not necessarily have physical contacts, groups, group identities and group norms can still be formed through social interaction. Once the social identity is formed, anonymity among group members in a computer-mediated world does not impede, but mostly enhance group bonds and performances [37]. For example, Ducheneaut et al. [14] studied the group formation in a virtual game world, and concluded that players usually remained anonymous to each other and often formed groups because of their complementary “skills”, which could help them tackle difficulties together in the game. Wessener and Pfister [52] argued that, for better group performance, computer-mediated system should support three phases of group formation, namely (1) the initiation (e.g., giving prompts to form groups or assigning groups in advance); (2) identification (e.g., finding group members); and (3) negotiation (e.g., balancing the preferences of group members and the goals of the group). The virtual environmental design and the pre-assigned groups of our virtual workshop well considered these factors.

Moving from the physical world to VEs may change social activities: creating new forms of interaction (e.g., use controllers to teleport in VEs), and new social norms to keep the virtual world in order [33]. Ackerman et al. [1] studied social regulation in an online game. They show specific social norms, defined to regulate the game world interactions and technical interventions, are imposed in game to automatically prevent or punish unwanted behaviours. Yee et al. [56] indicated that social interactions in VEs are governed by the same social norms (i.e., social norms of gender, interpersonal distance, and eye gaze) as social interactions in the physical world. Social VR has the potential to afford more social interaction than video conferencing, such as the ability to organically break off into small groups, or interacting with virtual objects in the scene [35]. Many commercial platforms have implemented novel social mechanics to stimulate social activities, such as designing a VE to simulate group discussion atmosphere, implementing built-in tools to enable users to stay in VEs and focus on the social tasks, or enabling users to use simple hand gestures to stop harassment [30].

Avatar realism and co-presence experiences have been recognised as two important factors for social VR activities. A recent study found that a realistic VE created more presence, and a cartoon avatar created stronger co-presence but was less trustworthy than a realistic one [25]. Steed and Schroeder [46] indicated that realistic avatar gaze may be important for one-on-one conversations in VR, but avatar distinctiveness is probably more important than realism when it comes to collaborating in large groups. Users often adjust their behaviour to compensate for the lack of social cues. For example, Roth et al. found that the absence of eye contact and facial expressions shifts the user's attention to other the tone of voice, and does not impede the task's execution [39]. The compensations also include exaggerating movements (e.g., flying up) when a user does not think their collaborator can see them [35].

3 RESEARCH QUESTIONS AND VIRTUAL WORKSHOP DESIGN

With the pandemic of 2020, numerous conferences and workshops were cancelled or rapidly adapted for remote attendance. Our workshop design was planned as a hybrid with a virtual VR component and in person component, but we pivoted to a fully virtual workshop. While redesigning our programme for virtual delivery, we explored three research questions: (RQ 1) Are metaphors from proxemics in the physical world useful for analysing proxemics in a virtual world? (RQ 2) What is the relationship between groups, environment and activity in a virtual workshop? (RQ 3) What social cues or environmental artefacts help or hinder the virtual workshop experience? Using data from our instrumented VE platform, observation notes, and semi-structured interviews, we measured user proximity, behaviours, and social experiences during the workshop. Our focus was to facilitate small group discussions, networking, addressing open-ended problems, and building communities in focused research areas.

3.1 On-Boarding Sessions

Prior to the workshop, we held on-boarding sessions to give participants a chance to test their technical setup, ask questions about the event, and familiarise themselves with Mozilla Hubs. This was particularly important for those not familiar with WASD keyboard movements. WASD is a common design for movement in video games [32] which uses the inverted T shape of those QWERTY keyboard keys for forward back left and right movements which are coupled with a mouse for a secondary action. The on-boarding also introduced many participants to the VE notion of flying, where a user can move into the sky and stay stationary. We also support participants during the opening the workshop with a warm up session to address any last minute technical issues.

3.2 Workshop Programme

The workshop officially started with a keynote and question session lasting forty minutes. Following the keynote was a ten minute break. In the next session, each participant gave a two minute pitch to introduce themselves in a session lasting fifty minutes. After a thirty minute break, the participants, broken into groups, began discussions in small breakout rooms accessed using portal linked from the main room. These discussions lasted one hour with two breaks. After the second break, the groups reconvened in the main room to present their discussions and close the workshop in a session lasting twenty minutes. The generous breaks throughout the day allowed participants to rest from in browser or VR headset, and provided social interaction interludes akin to "hallway conversations".

3.3 Workshop Environments

When designing a virtual world for the workshop, we opted for an outdoor lecture hall provided as a standard Mozilla Hubs space. The VE asset can be edited in a world builder from Mozilla called Spoke. This main room, called *Outdoor*




Fig. 2. *Outdoor Meetup* was the larger meeting space where keynotes and informal breaks took place.

Meetup, is a large open space with an equivalent area of seventy by forty meters. The standard scene has a large display screen in the space, a smaller screen in a side mini-amphitheatre, and several pre-populated tables. See Figure 2. We removed the tables as they only block movement in the space and have no virtual affordances. In the corners of the amphitheatre space, we added some instructions for using Mozilla Hubs as well as some Hawaiian themed items since the cancelled real-world event was to be in Honolulu. This room had some features disabled to keep participants from cluttering the space with drawings or objects.

There were three breakout rooms linked to the space with portals. Each breakout room was identical based on another standard Mozilla Hubs scene called “Lake Office”. Each of these rooms had wall space for screen-sharing and the room locks turned off so the breakout sessions could create objects, drawings, or other artefacts freely. The Lake Offices are equivalent to twenty by thirty meters (see figure 3). These breakout rooms had a small balcony overlooking a shallow lake providing some indoor and outdoor separation.

3.4 Informed Consent

To run this workshop while collecting data for our research, we asked each workshop participant for informed consent before the workshop. This consent included video recordings and photographs in the VE, observational notes, and instrumented tracking. This consent excluded recording chat or logging audio (outside of the livestreamed keynote on YouTube) as part of our privacy protections during the event. We explained steps we would take to anonymise the data, including removing participant names that would appear within the VE. Participants had the option to opt out of the semi-structured interviews.

4 INSTRUMENTATION, DATA, AND INTERVIEWS

Mozilla Hubs is an experimental, VR-friendly platform from Mozilla Mixed Reality. Mozilla Hubs works with immersive head mounted displays like Oculus Quest and traditional web browsers on a computer or mobile device like an iPhone. To collect data from our workshop we modified the Mozilla Hubs Client to track and log positional data, recorded observations during the workshop, and conducted semi-structured interviews after the workshop. The code described in this section for Mozilla Hubs Cloud data collection, the collected dataset, and scripts for analysis is openly available (see Appendix A).




Fig. 3. *Lake Office* was a smaller meeting space, where small group discussions took place. The three *Lake Office* rooms have access portals, as shown in the south-west corner of the *Outdoor Meetup* space in Figure 2.

4.1 Instrumented Virtual Environment

The Mozilla Hubs platform itself is entirely open source which allows people to stand up their own server instance with a custom client (called Hubs Cloud) on Amazon Web Services or Digital Ocean. We used a Hubs Cloud instance to run the workshop with instrumentation to facilitate the data collection. There is a Janus WebRTC Server for keeping track of all the messages, objects, and users in each room in Mozilla Hubs. To track where every user is, one could write code to constantly query the Janus server. However, this would create added load on an already stressed server that reaches its maximum capacity at thirty users. It is possible to modify and stand up a custom Janus server but that is generally not advised in the documentation. For our solution, we turned to client side logging.

Mozilla Hubs uses a two core web based frameworks, A-Frame and three.js, on the client. A-Frame is a web and HTML framework for virtual reality (WebVR); three.js is a cross-browser JavaScript library for creating and displaying animated 3D. These two provide the rendering engine that runs Hubs. More importantly, one can subscribe to tick events in the browser and on each tick log the needed data: (i) A unique identifier of the player (ii) Current UTC timestamp (iii) Have they entered in the room or are in the lobby? (iv) Where are they in the room (position vector as x, y, z)? (v) Where are they looking/facing in the room (direction vector, orientation quaternion)? (vi) What is the client's current rendered Frames-per-second? (vii) Are they muted? (viii) Are they talking (and if so, how loud on their input mic)? (ix) Do they have the spatial audio dampened? The overall framework allows one to add other listeners for events like use of the laser pointer or drawing pen.

Next, a HTTP server collected data from each client's POST data. However, each client logs data on every tick; this collection would add up to tens of tick POSTs per second. This would cause socket failures on the client and server. Here an open web-socket would allow data to flow freely. Instead, we opted for batch collecting ticks on the client then doing a bulk POST operation. We empirically tested the server and set a threshold of 4000 ticks per post. The collection server was a simple nginx HTTPd server with a python uWSGI application which logged JSON data via concatenation to disk.

4.2 Proxemic Dataset

Overall we collected 17,779 user input events (which includes the pose-position vectors) and 13,928 user info events (which track device usage). For this event, the logger captured positions every 1000 frames rendered. As the participants

were on devices of different capabilities, their effective frame-rates varied ($\mu = 43.67, \sigma = 15.3$) resulting in some participants having logged more frames than others. To address this, the pose-position data was resampled to ten frames per minute for standardisation. The resampled dataset had 37,457 user position logs organised into 6,619 frames from 26 users across four virtual spaces. Although the workshop had 27 registered attendees, one participant experienced technical trouble and could not attend the event.

The resampling process organised logs into frames, which are mapped to timestamps in UTC. For each frame, we calculated pairwise distance (euclidean in 3D space) and angle (based on position coordinates and direction vector) between each person. This is represented as two dictionaries of matrices, indexed by frame.

4.3 Observations

The workshop lasted 4.5 hours on April 29, 2020, which started at 16:10 (GMT+2) with co-organiser's introduction, and ended at 20:40 (GMT+2) with a virtual group photo. All the planned workshop activities, including the keynote speech, individual pitches, social breaks, group discussions and final presentations went smoothly. Some technical problems did exist, especially during the first half an hour, when a few participants were frequently reporting frozen frames and interrupting audio. However, these problems were solvable by refreshing and rejoin. Interestingly, even during social breaks, many users (14–18 users) chose to stay in the virtual world and network with each other. No participants reported technical problems during the group discussion in the breakout rooms. The participants in breakout room A spent their social breaks by jumping together into the virtual ocean.

4.4 Semi-structured Interviews

After the workshop, we conducted semi-structured interviews with nine participants (P1-P9, 6 females and 3 males). We aimed to collect participant feedback on overall experience such as quality of interaction/communication, their evaluation of Hubs and their thoughts of how the workshop compared to others they might have attended in the real world. Each interview was forty-five minutes long. In order to guide the conversation, we organised each interview into three parts. First was an opening question to encourage participants to talk about their overall experiences with prompts to remind them of specific sessions in the workshop to talk about for instance 'the keynote presentation' session or 'the breakout session'. Next, there was a discussion of the participants experience in the virtual workshop (in comparison to physical workshops if relevant), what delighted them about it, their frustrations about what did not work, the ease with which they partook in conversations, their ability to express themselves and their ability to form impressions of others. Finally, there were some user experience style questions asking for feedback about Hubs, in comparison to other video conference platforms, w.r.t audio quality, visual representation of the space, creation of avatar, usage of chat/emojis, and water cooler moments (if any).

5 RESULTS

Our results bring together analysis of the 26 users logged in the proxemic dataset (3 out of 26 users used HMDs), field observations, and 9 individuals given post-workshop interviews. Of the interview participants, eight attended the full workshop. P4 had significant technical difficulties during the workshop, but attended the workshop on-boarding session. P6 and P7 used an HMD (i.e., Oculus Quest) for the workshop, the other six participants all used web browsers (i.e., Chrome and Firefox).

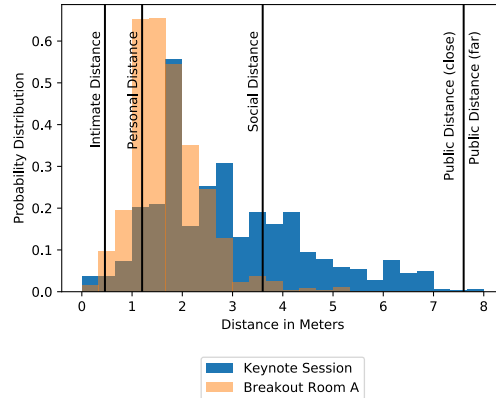


Fig. 4. This histogram shows the probability distribution of the nearest participant (euclidean distance in meters) in the Keynote Session compared to Breakout Room A. The proxemic zones [19] and labelled across the X axis. In the small group discussions in Breakout Room A, participants stood closer. There was a greater probability of standing with others in the personal and social distances as compared to the Keynote session. Collisions in the intimate distance occurred infrequently in both sessions and participants adjusted their social spacing when intimate distances occurred.

5.1 Social Cues and Proximity in Virtual Space

Using Hall’s proxemic zones [19], we analysed participants’ proximity and personal space in the VE. Our goal was to explore whether these zones were a useful construct for understanding how interaction unfolded during different workshop activities. Although previous research has shown the utility of measuring proxemic zones in VR, our dataset provides a continuous tracking at a high resolution to enable new kinds of analysis.

Figure 4 shows the probability distribution of collisions with the next nearest participant in the “intimate zone” (less than 0.5 meters) during the keynote session and breakout discussion in Room A. Mozilla Hubs does not prevent collisions or enforce “personal space” in software, but collisions in this space are still relatively rare. Such collisions can easily occur on “spawn” points when users enter the room at the same coordinates, but participants quickly moved away from others. Although participants discussed the importance of personal space, they also discussed some challenges. For example, some participants pointed out that it was not easy to perceive the positions and distance, especially to sense the things behind their avatar (P2, P3,P8).

We observed different behaviours across the proxemic zones during different workshop activities. Figure 4 shows the probability distribution of the next nearest participant during the Keynote Session and Breakout Room A. In Breakout Room A, participants formed cohesive small groups, where the next nearest participant was within the *personal* or *social* distances. During the Keynote session, the participants formed less cohesive groups, where the next nearest participant extended further away into the *public* zone. These insights demonstrate that impacts from activities or environments could be observed in the proxemic data.

Mozilla Hubs provides spatial and audio cues based on each person’s position, field of view, hand pose (if using an HMD), and microphone. Participants described the challenges of interpreting others’ intentions or starting conversations without additional cues like facial expressions and eye contact. P9 stated that “*I didn’t know when to cut into people’s talk, to express my opinion.*”. Similarly, P3 stated that “*I wasn’t sure whether people were paying attention. I would describe it as the biggest challenge in social VR where you don’t have facial cues like eye direction.*” Audio cues were the most




Fig. 5. Position data overview for the duration of the workshop in the *Outdoor Meetup* space. Activity took place most often in the centre of the room, but the floor decals in the space were an unexpected deterrent. Participants also avoided the large amphitheatre and no activities made specific use of this feature.

problematic during the workshop. All the participants mentioned that they couldn't have private chat or were uncertain how far away others could hear them. Audio often overlapped, and participants didn't want to interrupt others. If they turned on the spatial audio and moved away from others to chat, it could be hard to hear the workshop organisers and rejoin the main group. As P4 said, *"You just move from the main group and then you could have a private conversation. I know it is possible, but I didn't do it. In this [the social VR] setup, I was really focused on staying close to the group, not missing anything important."*

HMDs users benefited from additional social signals with tracked head and hand movements. Participants noticed how much more expressive participants in HMDS could be (P1–P3, P6–P8). P7, who wore an HMD, stated that *"wearing an HMD with moving hands made us more expressive. We seemed more important, because we can give more information and attract more attention using hand gestures."* Three participants (P2, P6, P7) wearing HMDs discussed how they perceived personal space and felt they could better pick up on social cues. P2 commented, *"I was not entirely aware [of the interpersonal space] because I couldn't easily turn my head to look behind me. So, I wouldn't have known if my avatar was standing right in front of somebody's eyesight. But I was conscious and always trying not to invade anyone's personal space."* All the participants mentioned that, it was much easier to sense the social cues and to have conversations in the breakout rooms. There was less audio overlapping, fewer people and smaller space where people cannot *fly* high. Group members were standing in a circle, which was easy to immediately identify the speaker. The circle felt like a realistic spacial relationship as in a real workshop. As P6 said, *"In the small room, it was more comfortable to talk because we were standing in a small circle. You could just tell people's intention from the small movements of the avatars. It was comfortable to interact with people."*

5.2 Groups and Social Interactions

Figure 5 shows all the positional data for the duration of the workshop in the *Outdoor Meetup* space. This was a large open space with activity focused in the centre of the area. The floor decals, visible as the hexagonal grid in Figure 5, appear to have discouraged participants from standing in those areas. Pedestrian traffic is incredibly sensitive to texture and appearance in physical settings [53], and VEs create similar experiences. The amphitheatre in the space was also poorly utilised, although we did not organise any activities that made use of it. Given the limited interactivity and resource provided by the static amphitheatre, this is in line with previous research [31].




Fig. 6. Heatmap of the *Outdoor Meetup* space during the Keynote Session overlaid with a quiver plot visualisation participants' field of view.




Fig. 7. Histogram of z-axis in *Outdoor Meetup* during the keynote and the second break. With Fly-mode enabled, participants could use the vertical space to gain a better view, but during breaks participants were more likely to constrain themselves to the floor plane.

During the keynote session, the speaker displayed their slides in the centre of the southern wall in *Outdoor Meetup*, opposite from the large display embedded in the environment. Figure 6 shows a heatmap of the participants' positions and a quiver plot of their field of view during the keynote presentation. During the keynote, participants kept the speaker within their field of view 54% of the time and maintained this as a shared point of attention. With plenty of room to spread out in the seventy by forty meter space, participants still occupied a relatively small area, 20×15 meters of the 70×40 room with sixteen people. Participants used the z-axis “fly mode” to simulate graded stadium seating while still maintaining personal space, as shown in Figure 7 where we measure proximity using 3 dimensions. The discussion breakout rooms were smaller with low ceilings, as seen in Figure 3. This resulted in closer distances amongst the participants (Figure 4) and limited use of flying within the space (Figure 7). In Mozilla Hubs, participants could fly through the ceiling and still fully participate in discussion, but this kind of behaviour was not observed. Intimate space collisions were also infrequent in the breakout rooms, as shown in Figure 4.

The workshop included four breaks to reduce fatigue and give participants time away from their screens. However, many participants stayed in the virtual space for informal conversations and networking. Figure 8 shows how participants formed groups during the first and second break sessions in the *Outdoor Meetup* space. At the first break, participants




Fig. 8. Heatmaps of the *Outdoor Meetup* space during scheduled breaks show where participants formed groups. During the first break, there was only one large group. During the second group, small groups surrounded a large group.



Fig. 9. Workshop participants during a small group breakout. The avatar second from the right has hand tracking enabled. Avatars without hand tracking have no hands rendered in the virtual environment.

formed a larger group that didn't have clear boundaries. During the second break, there was a large group surrounded by smaller satellite groups that became distinct points. While on break, participants were also more likely to remain on the floor plane, as shown in Figure 7. As participants became accustomed to the controls, using the virtual space, and sounds, they could form meaningful smaller groups. For example, participants figured out ways to attract attention or show that they were paying attention, such as *flying* to the visible positions to gain attention (P1–P9) or turning their avatar to face a speaker to show attention (P1–P9).

The participants discussed how their virtual workshop experience compared with face-to-face social interactions. When networking, participants found that the avatar faces the VE made it difficult to associate the avatar with the actual person. Participants felt there was additional effort to keep in contact with the people that they met during the workshop (P3, P5, P6), and they had fewer opportunities to encounter people by chance and start a conversation (P1, P2, P7). As P3 said, “*If we meet in person, and later we meet again, I would feel comfortable talking them, versus if we meet in avatar forms, next time when we actually see each other in person, we would still feel a need to reintroduce ourselves.*” Although we did not specifically address avatar capability or customisation in this research, participants discussed their experience with avatars during the workshop. The avatars in Mozilla Hubs are typically cartoonish or robotic representations, as shown in Figure 9. Most participants spent time selecting their avatar from Mozilla Hubs's defaults or customising using a third party tool⁴. The avatars are not photo-realistic, but the customisation allowed participants to represent themselves with a chosen appearance. Some participants pointed out that they prefer their avatars to be

⁴Avatar-Customizer used by the IEEEVR 2020 VR Conference, retrieved on Aug. 28, 2020.

realistic with expressive body language. They felt this would help them network and remember others more easily (P5, P7, P8). Other participants felt relaxed about their virtual appearance, didn't worry about others' judgement, and felt de-identified but well-represented by their avatars (P3–P6, P9). P4 stated that *"somehow having the characteristics related to myself, but not myself."* The avatar worked as an "equalizer (P6)", which enabled participants to be less anxious to approach senior people.

5.3 Play and Finding Familiarity

Participants used the time between activities to network and play creatively in the virtual space, for example jumping into the virtual ocean (P2, P6, P7-P9) and taking virtual selfies (P2, P8, P9). These playful activities became an effective icebreaker, for example P9 stated that *"before we jumped into the virtual ocean, everyone was talking and behaving professionally, a bit stiff. You hardly connect with people. But this [jumping into the virtual ocean] enables you to connect with everyone."* Participants also mentioned that the selfie feature was useful and enabled them to keep virtual memories, feel like they had been in a new place, and bond with the other participants (P2, P3, P5, P6, P8). P2 stated that *"taking selfies in social VR is definitely more interesting than a screenshot in video conferencing. You feel like you are actually in a new place, which is a big deal when you are in quarantine for a while (P2)."* Mozilla Hubs was new for most of the participants. This added difficulties for them in navigating the environments and utilising all the available features. The novelty of the VE could be distracting, for example P7 stated that *"the theme of the workshop is social VR. So, the Hubs itself became a conversation trigger, which is good. If it is a workshop about another topic, I am not sure the social VR platform would be more distracting than helpful."* Participants also mentioned difficulties such as saving or removing objects and selfies (P4, P7); feeling unnatural to turn head to see things (P3, P7); and difficulties in navigating to exact positions (P1, P3, P4).

The on-boarding session before the workshop was helpful for the participants to learn basic controls and social mechanisms in social VR (P1, P3, P8). As P3 mentioned, *"I was glad, throughout the training session, I had an opportunity to practice different controls. I wasn't sure how to figure out which people was taking, and later on, I learned this kind of social cues."* All participants mentioned that, when first entering the space, they moved around and explored the virtual space. They usually stood still when the speakers started to talk. *"When I entered, it was empty, so I moved around to just explore the virtual environment. When there were things happening, I mostly just stood at one point to have a good view (P6)."* Novelty also meant that participants weren't always sure what they were "allowed" to do or how they should behave. For example, P1 stated that *"the whole [social VR] experience at the beginning felt almost like my first conference experience. I was not really sure what I am doing, where I am, how to behave, how the space works, how loud my audio is and so on. The training session prepared us technically, but not psychologically (P1)."* Although the VE resembled face-to-face interactions in many ways, the social cues were not always enough to establish social norms. For example, physical proximity and personal space were easily managed using visual cues. Speech, however, was harder to control as the sound falloff did not always behave as it would in a face-to-face interaction.

5.4 Limitations

The workshop was a live event on the topic of Social VR; many of the participants were, of course, interested in the domain. However, 7 had no VR experience, 7 reported using VR once or twice briefly, 3 said they used VR several times, and 1 person stated they were a VR expert. As everyone was curious about VR and there to learn more, this shapes some of our qualitative findings. There were some technology failures and unforeseen limitations. We completed an on-boarding session to help workshop participants familiarise themselves with Mozilla Hubs and test their hardware

before the workshop. Many participants still experienced challenges, such as frequently dropping out (P1–P4) and audio issues (P5, P6, P8) likely attributed to work-from-home owners and smaller home Internet bandwidth connections. At the time of our study, Mozilla Hubs recommended a maximum of twenty-five participants per room for optimum performance⁵. This is variable based on client setup, but some of the participants experienced technical issues even with fifteen participants in the room. We also chose not to complete controlled comparisons, for example comparing equal numbers of browser and HMD users, during this live event. Our results are based on the unconstrained behaviours of the workshop participants, but future work could use our instrumented approach to complete more controlled experiments based on high precision proxemic logs. It would also be interesting to add audio signals to our analysis of proxemic logs to investigate the variations in spatial behaviour in response to aural differences. Although we subsequently added microphone volume tracking to our instrumentation code, this was not available during the workshop.

6 DISCUSSION

Open source VEs allow for unique instrumentation opportunities that can help advance research on social interaction in VEs. This paper reports an initial step in that direction, which we hope will support future work and enable new kinds of evaluations. In the following, we will discuss relevant challenges and surprising findings from our workshop, including the on-boarding process, the importance of the environment design, the influence of users' hardware, and the novel interactions made possible by virtual environments. We conclude with some design insights to inspire future work.

Novel experiences always face on-boarding challenges. These challenges typically include usability, sociability, and learning curve components. The latter tends to be a product of the first two. During the interviews, participants discussed how the virtual workshop was similar to a “first conference” experience. In this case, the lack of familiarity with the environment might have been an advantage. Many participants had their first experiences with Mozilla Hubs during our on-boarding process, which established the social norms of an academic workshop. VEs are frequently used for events like DJ parties or art museum experiences, but one of our main objectives was to create a conference experience without feeling game-like or drawing from gaming metaphors. For participants familiar with game platforms like Minecraft or World of Warcraft, developing new norms and standards for non-game experiences could've been an additional challenge. The on-boarding process was positively received and like any other novel interface, it was a vital part of the virtual workshop experience. It's advisable to have both a written short introduction to the hosting platform and a short pre-workshop session on the platform in order to familiarise participants with the environment and expected social norms [27].

Mozilla Hubs currently does not provide much *social translucence* [15], which is most problematic when moving between different linked environments. Social translucence, or more simply social awareness, states digital system design should support coherent behaviour by making actions visible to all participants. In a simple example, social cues or other observed behaviours should be communicated—like User B left the room or User Z is typing—to maintain visibility, awareness, and accountability. While our environment contained portals linking the breakout rooms to the main room, there was no visibility *between* rooms. For example, the only way to see if someone is in a different room is to travel there [27]. This can be a disruptive experience, which requires one to fully exit one space before gaining awareness of the state of the new space. For observers, someone entering a portal instantly disappears from the current room as if they lost connectivity or closed their browser window. These portals are uni-directional and provide limited

⁵What is the capacity of a Hubs room?, accessed September 2020.

information about the destination, for example if one will be able to easily return if they use the portal. Increased social translucence must provide feedback that decreases the disruption of moving between spaces, gives awareness of the state of destination, and makes movements between spaces visible.

In guided virtual workshops, such as ours, translucency could provide a way for organisers to ensure all attendants are stationed in the rooms they are meant to be at any given time. In their paper, Lu et al. [27] talk of different motivations to attend a conference. Ultimately, all attendees want to learn about new research findings and connect with others. However, remote attendees wanted to connect with new people while local attendees were interested in socialising with existing or potential collaborators they already knew about. Initialising a conversation is made more difficult especially if you are a remote attendee with no easily recognisable contacts as implied by some participants (P3, P9). In the future, it would be interesting to investigate how to allow attendees and moderators to have a better awareness of who else is attending, which virtual room they are in and also enable them to capture, interpret and express the right social cue to turn-take and indicate interest in joining a group conversation regardless of which device the attendee is using. In our workshop, we saw some unintended but delightful interactions and behaviours displayed around the ‘tiki corner’ and the ‘ocean’ outside the breakout room. It appears that even in serious fit-for-purpose virtual space, there is some mileage in including what may seem as superfluous objects in order to encourage ice-breaking social behaviour—jumping into the ocean with strangers or taking group selfies with friends.

Based on our quantitative proxemic dataset, we learned about the influence that a VE can have on the participants and their behaviours. The *Outdoor Meetup* space included a large screen, a podium and a mini-theatre, intended for presentations. The keynote speaker decided instead to project his slides to the back of the room, making the reserved keynote space irrelevant. The ability to dynamically reconfigure space easily in a VE allows for unexpected space occupancy and group formations. Dynamic changes like this can change the way people use virtual space in way unanticipated by the environment’s designers. While some designed features proved irrelevant, some incidental features had unexpected impacts. For example, the floor pattern rendered as a texture on the floor plane had a strong influence on the positioning of the participants, as shown in Figure 5. Participants avoided the floor texture and appeared to surf the edge when moving in the space. These insights open new areas of research around data-driven design of VEs. For example, can we gather and process large-scale proxemic and movement data in order to automatically generate spaces for different purposes and influence pedestrian flow and social grouping? Can the environments self-adapt based on the actual use in order to increase engagement? Mathematical approaches to space design [22] could have new relevance in VEs when proxemic data is continuously available.

As a 3D virtual environment, spatial relationships form a large part of the user experience in Mozilla Hubs. The interface provides a field of view, and users must navigate the VE and interact with others through this view. Our participants were aware of spatial relationships and spread out during the workshop. The amount of personal space participants maintained depended on the activity and the design of the environment. Although, we introduced participants to Mozilla Hubs and its functionality, there were some virtual skills we didn’t cover in that initial introduction. However, similar to in real world situations, participants could observe others’ actions, like flying, and would ask each other how they too might fly. Similar to other reports [27], flying was seen as a fun activity but it wasn’t always deemed “acceptable,” as we only observed flying regularly during the keynote. This ‘stadium seating pattern’ was not observed during other kinds of activities like small breakout groups. The *Lake Office* rooms may have created a sense of “perceived containment” from incorrectly assuming walls would act as barriers or a social expectation to remain visible within the walls. Flying in the larger *Outdoor Meetup* space was also rare during informal breaks. Social expectations may discourage such behaviour that could appear like aggressively seeking privacy.

Previous works [41, 43, 47] demonstrate that asymmetry between participants on different systems affects their collaboration in a task. Participants with a network disadvantage, for instance, were unlikely to emerge as the ‘leader’ in a collaborative task and participants using an HMD were more likely to emerge as the ‘leader.’ In our workshop, in order to handle this, we had assigned moderators who were the automatic guides regardless of device capability. Still, the experiences of users with HMDs and users without HMDs were significantly different. In Mozilla Hubs, participants using HMDs have hands rendered as part of their avatar, while participants using a web browser do not. This allowed the HMD users to be more expressive and gesture. Having hands in the VE aided casual conversation and allowed HMD users to add additional information to their speech, for example using diegetic gestures during their presentation. This was observed in other virtual events as well [27], where audiences watching a presentation wanted to applaud at the end. Although an imbalance of expression was noticeable, there were also some advantages when using an HMD. HMD users could not share a screen or a document as easily as the moderators in the fast paced breakout sessions. Finally, participants using a browser on a computer or a mobile device have a more standard ‘video call’ feeling as they could move the window aside while looking for a document or taking notes. The moderators in all breakout rooms chose to attend the workshop using their desktop in order to facilitate and take live notes. This allowed for a relatively long but very interactive workshop—half a working day. The lesson we learned regarding the influence of the local configuration is the need for adequate customisation or intervention elements to minimise their imbalanced experiences. In the future, it would be interesting to study groups of workshop attendees where the role of a moderator is undefined. For instance, does the note-taker and facilitator roles in a small group always fall upon the one attendee on a desktop if everyone else is attending using an HMD, even though at first glance it might seem as if the desktop is not the most advanced device to use when attending virtual workshops?

7 CONCLUSION

Remote events can benefit from virtual environments. In the context of an academic workshop, we found that proxemic interactions are congruent with those in the physical, although VEs are not constrained by physics and we observed that flying could add a new dimension to personal space. Our results demonstrate how group formations adapt to the world/room size and the ability to create conversation clusters. Further our approach can provide novel insights into the effect of the environment textures, shapes, and routing. While some social cues were analogous to real world interactions, there is a need for expression, indicators, and translucence when using a generic VE like Mozilla Hubs. Being open source, Mozilla Hubs is extensible to gain precise user actions and position data at tick resolution to analyse and answer questions on social behaviours as well as providing ample data for future insights and possible generative layouts to support collaborative meetings and tasks.

ACKNOWLEDGMENTS

The authors would like to thank to Yulius Tjahjadi for standing up our Mozilla Hubs Cloud beta instance in AWS and Blair MacIntyre at Georgia Tech for keynoting the workshop. We extend additional thanks to Liv, Elgin, and the rest of Mozilla Hubs team for their assistance in setting up the VE. We also thank all the workshop participants for joining us during this event and making this research possible.

This research was partially funded by the European Commission as part of the H2020 program under the grant agreement #762111 “VRTogether” <http://vrtogether.eu/> (accessed January 2021).

This research was partially funded by the European Commission as part of the H2020 ERC project ViAJeRo, grant agreement #835197. <https://viajero-project.org/> (accessed January 2021).

REFERENCES

- [1] Mark S. Ackerman, Jack Muramatsu, and David W. McDonald. 2010. Social Regulation in an Online Game: Uncovering the Problematics of Code. In *Proceedings of the 16th ACM International Conference on Supporting Group Work* (Sanibel Island, Florida, USA) (*GROUP '10*). Association for Computing Machinery, New York, NY, USA, 173–182. <https://doi.org/10.1145/1880071.1880101>
- [2] Michael Argyle. 1988. *Bodily Communication*. Routledge, Abingdon-on-Thames, England.
- [3] Michael Argyle and Janet Dean. 1965. Eye-Contact, Distance and Affiliation. *Sociometry* 28, 3 (1965), 289–304. <http://www.jstor.org/stable/2786027>
- [4] Michael Argyle and Peter Trower. 1979. *Person to person: ways of communicating*. Harper and Row, New York, NY, USA.
- [5] Jeremy N. Bailenson, Jim Blascovich, Andrew C. Beall, and Jack M. Loomis. 2001. Equilibrium Theory Revisited: Mutual Gaze and Personal Space in Virtual Environments. *Presence: Teleoper. Virtual Environ.* 10, 6 (Dec. 2001), 583–598. <https://doi.org/10.1162/105474601753272844>
- [6] Jeremy N. Bailenson, Jim Blascovich, Andrew C. Beall, and Jack M. Loomis. 2003. Interpersonal distance in immersive virtual environments. , 819–833 pages. <https://doi.org/10.1177/0146167203029007002>
- [7] John Bowers, James Pycock, and Jon O'Brien. 1996. Talk and Embodiment in Collaborative Virtual Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, British Columbia, Canada) (*CHI '96*). Association for Computing Machinery, New York, NY, USA, 58–65. <https://doi.org/10.1145/238386.238404>
- [8] Barry Brown. 2020. Notes on Running an Online Academic Conference or How We Got Zoombombed and Lived to Tell the Tale. *Interactions* 27, 4 (7 2020), 16–21. <https://doi.org/10.1145/3406108>
- [9] Barry Brown and Marek Bell. 2004. CSCW at Play: 'there' as a Collaborative Virtual Environment. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work* (Chicago, Illinois, USA) (*CSCW '04*). Association for Computing Machinery, New York, NY, USA, 350–359. <https://doi.org/10.1145/1031607.1031666>
- [10] A. Bönsch, S. Radke, H. Overath, L. M. Asché, J. Wendt, T. Vierjahn, U. Habel, and T. W. Kuhlen. 2018. Social VR: How Personal Space is Affected by Virtual Agents' Emotions. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, New York, NY, USA, 199–206. <https://doi.org/10.1109/VR.2018.8446480>
- [11] Laura Cabrera-Quiros, D. Tax, and Hayley Hung. 2020. Gestures In-The-Wild: Detecting Conversational Hand Gestures in Crowded Scenes Using a Multimodal Fusion of Bags of Video Trajectories and Body Worn Acceleration. *IEEE Transactions on Multimedia* 22 (2020), 138–147.
- [12] Marcus Carter, Fraser Allison, John Downs, and Martin Gibbs. 2015. Player Identity Dissonance and Voice Interaction in Games. In *Proceedings of the 2015 Annual Symposium on Computer-Human Interaction in Play* (London, United Kingdom) (*CHI PLAY '15*). Association for Computing Machinery, New York, NY, USA, 265–269. <https://doi.org/10.1145/2793107.2793144>
- [13] Alvin Chin, Bin Xu, Fangxi Yin, Xia Wang, Wei Wang, Xiaoguang Fan, Dezhi Hong, and Ying Wang. 2012. Using Proximity and Homophily to Connect Conference Attendees in a Mobile Social Network. In *Proceedings of the 2012 32nd International Conference on Distributed Computing Systems Workshops (ICDCSW '12)*. IEEE Computer Society, USA, 79–87. <https://doi.org/10.1109/ICDCSW.2012.56>
- [14] Nicolas Ducheneaut, Nicholas Yee, Eric Nickell, and Robert J. Moore. 2006. "Alone Together?": Exploring the Social Dynamics of Massively Multiplayer Online Games. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montréal, Québec, Canada) (*CHI '06*). Association for Computing Machinery, New York, NY, USA, 407–416. <https://doi.org/10.1145/1124772.1124834>
- [15] Thomas Erickson and Wendy A. Kellogg. 2000. Social Translucence: An Approach to Designing Systems That Support Social Processes. *ACM Trans. Comput.-Hum. Interact.* 7, 1 (March 2000), 59–83. <https://doi.org/10.1145/344949.345004>
- [16] Thomas Erickson, N. Sadat Shami, Wendy A. Kellogg, and David W. Levine. 2011. Synchronous Interaction among Hundreds: An Evaluation of a Conference in an Avatar-Based Virtual Environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (*CHI '11*). Association for Computing Machinery, New York, NY, USA, 503–512. <https://doi.org/10.1145/1978942.1979013>
- [17] Doron Friedman, Anthony Steed, and Mel Slater. 2007. Spatial social behavior in second life. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 4722 LNCS (2007), 252–263. https://doi.org/10.1007/978-3-540-74997-4_{ }23
- [18] B. Guo, C. Chen, D. Zhang, Z. Yu, and A. Chin. 2016. Mobile crowd sensing and computing: when participatory sensing meets participatory social media. *IEEE Communications Magazine* 54, 2 (2016), 131–137.
- [19] Edward Hall. 1969. *The Hidden Dimension : man's use of space in public and in private*. Anchor Books, Hamburg, Germany. 217 pages.
- [20] Leslie A. Hayduk. 1983. Personal space: Where we now stand. *Psychological Bulletin* 94, 2 (1983), 293–335. <https://doi.org/10.1037//0033-2909.94.2.293>
- [21] Heiko Hecht, Robin Welsch, Jana Viehoff, and Matthew R. Longo. 2019. The shape of personal space. *Acta Psychologica* 193, April 2018 (2019), 113–122. <https://doi.org/10.1016/j.actpsy.2018.12.009>
- [22] B. Hillier and J. Hanson. 1988. *The social logic of space*. Cambridge University Press, Cambridge, England.
- [23] Michael A. Hogg and John C. Turner. 1985. Interpersonal attraction, social identification and psychological group formation. *European Journal of Social Psychology* 15, 1 (1 1985), 51–66. <https://doi.org/10.1002/ejsp.2420150105>
- [24] Hayley Hung and Daniel Gatica-Perez. 2010. Estimating cohesion in small groups using audio-visual nonverbal behavior. *IEEE Transactions on Multimedia* 12, 6 (10 2010), 563–575. <https://doi.org/10.1109/TMM.2010.2055233>
- [25] Dongsik Jo, Ki-Hong Kim, and Gerard Jounghyun Kim. 2017. Effects of Avatar and Background Types on Users' Co-presence and Trust for Mixed Reality-Based Teleconference Systems. In *Proceedings of the International Conference on Computer Animation and Social Agents*. Wiley, New York, NY, USA, 27–36.

- [26] Taemie Kim, Agnes Chang, Lindsey Holland, and Alex Sandy Pentland. 2008. Meeting Mediator: Enhancing Group Collaboration using Sociometric Feedback. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work* (San Diego, CA, USA) (CSCW '08). Association for Computing Machinery, New York, NY, USA, 457–466. <https://doi.org/10.1145/1460563.1460636>
- [27] D. A. Le, B. MacIntyre, and J. Outlaw. 2020. Enhancing the Experience of Virtual Conferences in Social Virtual Environments. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, New York, NY, USA, 485–494. <https://doi.org/10.1109/VRW50115.2020.00101>
- [28] Jie Li, Yiping Kong, Thomas Röggl, Francesca De Simone, Swamy Ananthanarayan, Huib De Ridder, Abdallah El Ali, and Pablo Cesar. 2019. Measuring and understanding photo sharing experiences in social virtual reality. In *Conference on Human Factors in Computing Systems - Proceedings*. Association for Computing Machinery, New York, New York, USA, 1–14. <https://doi.org/10.1145/3290605.3300897>
- [29] Joan Llobera, Bernhard Spanlang, Giulio Ruffini, and Mel Slater. 2010. Proxemics with Multiple Dynamic Characters in an Immersive Virtual Environment. *ACM Trans. Appl. Percept.* 8, 1, Article 3 (Nov. 2010), 12 pages. <https://doi.org/10.1145/1857893.1857896>
- [30] Joshua McVeigh-Schultz, Anya Kolesnichenko, and Katherine Isbister. 2019. Shaping Pro-Social Interaction in VR: An Emerging Design Framework. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300794>
- [31] Robert Moore, E. Hankinson Gathman, and Nicolas Ducheneaut. 2009. From 3D space to third place: The social life of small virtual spaces. *Human Organization* 68, 2 (2009), 230–240. <https://doi.org/10.17730/humo.68.2.q673k16185u68v15>
- [32] Richard Moss. 2018. *The Secret History of Mac Gaming*. Unbound, London, UK, Chapter The making of Dark Castle, 121–135. https://www.gamasutra.com/view/news/315720/The_making_of_Dark_Castle_An_excerpt_from_The_Secret_History_of_Mac_Gaming.php
- [33] Fares Moustafa and Anthony Steed. 2018. A longitudinal study of small group interaction in social virtual reality. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology, VRST*, Vol. 10. Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/3281505.3281527>
- [34] Carman Neustaedter, Gina Venolia, Jason Procyk, and Daniel Hawkins. 2016. To Beam or Not to Beam: A Study of Remote Telepresence Attendance at an Academic Conference. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing* (San Francisco, California, USA) (CSCW '16). Association for Computing Machinery, New York, NY, USA, 418–431. <https://doi.org/10.1145/2818048.2819922>
- [35] Catlin Pidel and Philipp Ackermann. 2020. Collaboration in Virtual and Augmented Reality: A Systematic Overview. In *Augmented Reality, Virtual Reality, and Computer Graphics*. Springer, Cham, Berlin, Germany, 141–156. https://doi.org/10.1007/978-3-030-58465-8_10
- [36] T.T. Postmes, R. Spears, and M. Lea. 1999. Social identity, group norms, and "deindividuation" Lessons from computer-mediated communication for social influence in the group. , 164–183 pages.
- [37] Tom Postmes, Russell Spears, Khaled Sakhel, and Daphne de Groot. 2001. Social Influence in Computer-Mediated Communication: The Effects of Anonymity on Group Behavior. *Personality and Social Psychology Bulletin* 27, 10 (10 2001), 1243–1254. <https://doi.org/10.1177/01461672012710001>
- [38] Surangika Ranathunga, Stephen Cranefield, and Martin Purvis. 2012. Identifying events taking place in Second Life virtual environments. *Applied Artificial Intelligence* 26, 1-2 (2012), 137–181. <https://doi.org/10.1080/08839514.2012.629559> arXiv:<https://doi.org/10.1080/08839514.2012.629559>
- [39] D. Roth, J. Lugin, D. Galakhov, A. Hofmann, G. Bente, M. E. Latoschik, and A. Fuhrmann. 2016. Avatar realism and social interaction quality in virtual reality. In *2016 IEEE Virtual Reality (VR)*. IEEE, New York, NY, USA, 277–278. <https://doi.org/10.1109/VR.2016.7504761>
- [40] Albert E. Scheflen and Alice Scheflen. 1972. *Body Language and the Social Order: Communication As Behavioral Control*. Prentice Hall Direct, Des Moines, IA, USA.
- [41] Ralph Schroeder, Anthony Steed, Ann-Sophie Axelsson, Ilona Haldal, Åsa Abelin, Josef Widestrom, Alexander Nilsson, and Mel Slater. 2001. Collaborating in networked immersive spaces : as good as being there together? *Computers & Graphics* 25, 5 (2001), 781–788. Collaborating in networked immersive spaces: as good as being there together? xD;TY - JOUR Y2 - Oct ISI:000171668700006 484AU COMPUT GRAPH-UK xD:<http://www.equator.ac.uk/var/uploads/Schroeder2001.pdf>.
- [42] S. Shirmohammadi, S. Hu, W. T. Ooi, G. Schiele, and A. Wacker. 2012. Mixing virtual and physical participation: The future of conference attendance?. In *2012 IEEE International Workshop on Haptic Audio Visual Environments and Games (HAVE 2012) Proceedings*. IEEE, New York, NY, USA, 150–155. <https://doi.org/10.1109/HAVE.2012.6374455>
- [43] Mel Slater, Amela Sadagic, Martin Usoh, and Ralph Schroeder. 2000. Small-Group Behavior in a Virtual and Real Environment: A Comparative Study. *Presence* 9 (02 2000), 37–51. <https://doi.org/10.1162/105474600566600>
- [44] Mel Slater and Maria V. Sanchez-Vives. 2016. Enhancing Our Lives with Immersive Virtual Reality. *Frontiers in Robotics and AI* 3 (2016), 74. <https://doi.org/10.3389/frobt.2016.00074>
- [45] Mel Slater, Martin Usoh, and A. Steed. 1994. Depth of presence in virtual environments. *Presence: Teleoperators and virtual environment* 3, 2 (1994), 130–144. <https://doi.org/10.1007/s13398-014-0173-7.2>
- [46] Anthony Steed and Ralph Schroeder. 2015. Collaboration in immersive and non-immersive virtual environments. In *Immersed in Media: Telepresence Theory, Measurement and Technology*. Springer International Publishing, Berlin, Germany, 263–282. https://doi.org/10.1007/978-3-319-10190-3_11
- [47] A. Steed, M. Slater, A. Sadagic, A. Bullock, and J. Tromp. 1999. Leadership and collaboration in shared virtual environments. In *Proceedings IEEE Virtual Reality (Cat. No. 99CB36316)*. IEEE, New York, NY, USA, 112–115. <https://doi.org/10.1109/VR.1999.756941>
- [48] Duck Steve. 1998. *Human Relationships*. SAGE, London, England.
- [49] Kohske Takahashi, Tobias Meilinger, Katsumi Watanabe, and Heinrich H. Bühlhoff. 2013. Psychological influences on distance estimation in a virtual reality environment. *Frontiers in Human Neuroscience* 7, SEP (2013), 1–7. <https://doi.org/10.3389/fnhum.2013.00580>

- [50] Shari M. Trewin, Mark R. Laff, Anna Cavender, and Vicki L. Hanson. 2008. Accessibility in Virtual Worlds. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems* (Florence, Italy) (*CHI EA '08*). Association for Computing Machinery, New York, NY, USA, 2727–2732. <https://doi.org/10.1145/1358628.1358752>
- [51] Greg Wadley, Marcus Carter, and Martin Gibbs. 2015. Voice in virtual worlds: The design, use, and influence of voice chat in online play. *Human-Computer Interaction* 30, 3-4 (2015), 336–365. <https://doi.org/10.1080/07370024.2014.987346>
- [52] M. Wessner and H. R. Pfister. 2001. Group formation in computer-supported collaborative learning. In *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work*. Association for Computing Machinery (ACM), New York, New York, USA, 24–31. <https://doi.org/10.1145/500286.500293>
- [53] William H. Whyte. 1982. *The Social Life of Small Urban Spaces*. Project for Public Spaces, New York, NY, USA.
- [54] Laurie M. Wilcox, Samuel Elfassy, Cynthia Grelik, and Robert S. Allison. 2006. Personal Space in Virtual Reality. *ACM Transactions on Applied Perception* 3, 4 (2006), 412–428. <https://doi.org/10.1145/1190036.1190041>
- [55] Bob G. Witmer and Michael J. Singer. 1998. Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and Virtual Environments* 7, 3 (1998), 225–240. <https://doi.org/10.1162/105474698565686>
- [56] Nick Yee, Jeremy N. Bailenson, Mark Urbanek, Francis Chang, and Dan Merget. 2007. The unbearable likeness of being digital: The persistence of nonverbal social norms in online virtual environments. *Cyberpsychology and Behavior* 10, 1 (2 2007), 115–121. <https://doi.org/10.1089/cpb.2006.9984>

A SUPPLEMENTAL MATERIALS

The supplemental materials with this article include:

- JavaScript code to instrument Mozilla Hubs Cloud for data logging. Two of the detailed instrumented items, 4.1(viii) and 4.1(ix), were later additions to the collection code and not available during our study.
- The anonymous dataset collected from our workshop.
- A collection of Python Jupyter Notebooks to clean and analyse the data, generate the figures and animation in this paper.

These materials are available in the ACM Digital Library, as well as, on Github at <https://github.com/ayman/hubs-research-acm-chi-2021> (accessed January 2021).