

Situated Practices of Looking: Visual Practice in an Online World

Lilly Irani, Gillian R. Hayes, and Paul Dourish

Department of Informatics
Donald Bren School of Information and Computer Sciences
University of California, Irvine
Irvine, CA 92697-3440, USA
{lirani, gillianrh, jpd}@ics.uci.edu

ABSTRACT

Graphical virtual worlds are increasingly significant sites of collaborative interaction. Many argue that the simulation of the everyday environment makes them particularly effective for collaboration. Based on a study of visual practice in Second Life, we argue: first, that the practice of looking is more varied than it might at first seem; second, that we need to look beyond the virtual in understanding virtual worlds; and third, that implementations blend interactional practice. We suggest that the value of virtual worlds as sites of collaboration might lie more in their richness and openness to appropriation than in their simulation of everyday interaction.

Author Keywords

Looking, virtual worlds, collaborative virtual environments, situated practice

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Group and Organization Interfaces: Theory; H5.m. Information interfaces and presentation (e.g., HCI): User Interfaces: Graphical User Interfaces; H5.m. Information interfaces and presentation (e.g., HCI): Group and Organization Interfaces: Synchronous Interaction

INTRODUCTION

A persistent theme in technology-oriented CSCW research has been the metaphorical relationships between online and offline interactions. Real world interactions are frequently used as a model for interaction design [22], and face-to-face interactions are held up as a gold standard against which to judge new technologies [11]. Although some researchers have questioned “the real world baseline” as an evaluative rubric [7, 19], it remains a source of metaphors for how collaborative environments – and actions in them – are

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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW'08, November 8–12, 2008, San Diego, California, USA.
Copyright 2008 ACM 978-1-60558-007-4/08/11...\$5.00.

understood. Yet, while virtual worlds become comprehensible by drawing upon metaphors from daily life – ground, bodies, and walking, for example – they are also engaging precisely because they do not slavishly enforce the rules of the physical world. In virtual worlds, one can often fly, adopt a new persona, or carry on five conversations at once.

One critical aspect of virtual worlds’ comprehensibility is that actions become mutually comprehensible through their performance in a common visual space. Looking and visibility have been noted as important features of computer mediated experiences, such as media spaces [7, 11] and collaborative virtual environments (CVEs) [6, 16, 22].

Looking, of course, is only one part of the complex set of activities, such as gesturing and talking, interwoven into embodied action. We focus on looking not to argue its centrality, but instead to extend and qualify prior, fully-embodied understandings of the social life of virtual worlds. Unpacking activity from the focal point of looking provides a deeper understanding of activity within and around virtual worlds.

Discussions of looking, however, have typically been focused on gaze, either among avatars or among remote collaborators. Virtual worlds are most often discussed as a window into a simulated geography. However, the mundane realities of CVEs as applications, with their own menus, bugs, and obligations to operating system conventions, are also consequential to the practices of looking in virtual worlds. It is this situated, holistic analysis of looking in and around CVEs that we take up here.

To understand how these issues unfold in everyday computer-mediated communication (CMC) we studied two communities in Second Life as interviewers and participant-observers. Each community focused on the experience of a particular minority culture.

The title of this paper is taken from a comment of Harold Garfinkel’s, in *Studies in Ethnomethodology* [10]: “When I speak of accountable [as a property of social settings] ... I mean observable-and-reportable, *i.e.*, available to members as situated practices of looking and telling.” Garfinkel’s

concern is with how socially competent practice consists in part of the ability to see the world as organized in appropriate terms; to look at the world as a bodyguard, for example, is to see it filled with threats and escape routes.

In this paper, we are concerned specifically with the situated practices of looking, how looking is accomplished in a virtual world, and the associated questions of mutual interpretability of those acts among people in that world. We argue, first, that “looking” is a considerably more varied practice in the virtual world than one might imagine, manifesting itself in many different ways according to the needs of the moment; second, that to understand visual practice within a virtual world, we must examine it more broadly, as an interface that appears on a computer and is controlled with keyboard and mouse, rather than focusing solely on the simulation of action in virtual geographies; and, third, that if we are truly to look at these as situated practices, then the constraints and specificities of software implementations, such as bugs that deviate from the design, must be considered.

After a brief review of related work, we discuss data from our ongoing study in Second Life. We discuss these results in terms of studies of the socially situated nature of visual practice and conclude by exploring the consequences of this analysis for the study of virtual environments.

RELATED WORK

Our research builds on, extends, and qualifies previous work on looking and vision, both in the physical world and online in collaborative virtual environments (CVEs) and media spaces. In this section, we focus on the related work in CVEs and media spaces, including studies of long-term use of these mediated spaces.

Much work on looking in mediated spaces investigates the role of looking and gaze in achieving social order and communication. Much research has explored gaze and looking as social ordering resources in CVEs [2, 6, 14-16]. Communication in CVEs face several challenges related to looking and gaze: distorted or unavailable visual cues disrupt turn-taking and speaker influence [14], reduced field of vision makes visual access to collaborators more difficult [15], reduced gaze and pointing cues make it difficult to assess what others in the virtual world are looking at [15], and automatically generated gaze often generates misunderstandings [2, 22]. Most of these works propose that CVEs can improve by working towards better supporting the sorts of gaze and gesture cues available in the real or physical world.

Another line of research on looking investigates affective properties of virtual gaze, drawing from social psychological perspectives. First, there is evidence that the effects of gaze on the proxemics of avatars echoes how people might space themselves in the physical world. Further, these similarities are strengthened if a user believes the avatar is controlled by a person [3] rather than

algorithmically. In person-to-person audio conference mediated by avatars, avatar gaze inferred from the conversation also has been shown to have a more positive effect on feelings of presence than random avatar gaze [9]. This line of research, often conducted as laboratory experiments, has focused typically on comparing the effects of avatar behavior to that of actual people or video representations, implicitly drawing on the “real world baseline” in evaluating avatars.

To supplement short-term engagements in virtual worlds, as studied in the lab, longer-term investigations of media spaces and virtual worlds shed light on how new practices can emerge around media space affordances. For example, Dourish *et al.* described how people developed new practices of gazing into media space cameras as they came to understand how video mediated and transformed their image to their audience [7]. Little is known, however, about the development of embodied practices among long-term CVE communities, and how those practices might be specific to interface particularities – both features and quirks – of that CVE. Our research addresses this gap.

Research has also tended to focus on the virtual world as the interaction of avatars within a virtual geography. We have less knowledge about how the experience of a virtual world space is shaped by traditional aspects of an application, such as windowing mechanisms and menus, and how the application interacts with other applications. When Hindmarsh, for example, proposed making virtual world looking accountable by highlighting target objects in the space, targets such as profiles, menus, and maps were implicitly excluded [16]. More recent analyses have mentioned non-simulation aspects of the interface. Hindmarsh noted that users’ physical monitor size has important consequences for mutual visibility in CVEs [1]. Brown’s analysis of *There* [6] and Moore *et al.*’s analyses of *World of Warcraft* and *There* [22] also note consequences of certain interface design and implementation artifacts on CVE experience. While these projects mention non-simulation aspects of the interface as exceptions to the simulation experience, our holistic analysis takes simulation and non-simulation aspects of the experience as first-order phenomena of interest.

METHOD

The research reported in this paper grew out of a continuing interest in how sharing, trust, intimacy, and community emerge out of everyday social activity in online communities. We have been investigating these issues in the online world *Second Life* (SL). Launched in 2003, *Second Life* (www.secondlife.com) is an open-ended, 3D collaborative virtual environment. Like *World of Warcraft* (WoW), *Active Worlds*, and *There* [6], SL users move through a persistent world using a customizable avatar that allows them to interact with in-world objects. Objects and actions in the world are available not only through the user interface, but also through an API that allows them to create

objects with custom behaviors and provides access to world states not otherwise available through the user interface. SL differs from massively multiplayer online roleplaying games, such as WoW, in that it has neither predetermined narratives nor game objectives. Similar to *There*, SL includes a two-dimensional (2D) web information space containing maps of the virtual world, calendars of events, and the ability to launch SL locations from the Web Browser. SL hosts several long-standing communities that provide support, information, and advocacy for people with disabilities and medical conditions.

In SL, we studied two communities: a disability activist community and an informal genderqueer¹ social group. These groups were of particular interest for four reasons. First, community members participate in SL for support, information, and sociality, rather than technological enthusiasm. Second, the first group is a long-standing community (it has existed for a decade through IRC and the web and three years in SL). As such, it has evolved relatively stable community practices. The other group provided a case study at the other edge, as a relatively informal and new community. Third, as marginalized populations, members of these communities may be particularly conscious of and attentive to issues surrounding privacy and intimacy [8]. Fourth, disability and minority activism can often be crucially concerned with identity politics and self-presentation [12]. In addition to these communities, our study also included interactions and conversations with a broader set of users encountered through snowball sampling and virtual world proximity.

The challenges of articulating practices of looking and visibility naturally drew us to a more observation-focused inquiry. Our approach draws from virtual ethnography [18], which provides a method for analyzing digital communicative artifacts, and multi-sited – across SL, web forums, and the “real world” – ethnography [21]. We completed 60 hours of in-world participant-observation and 18 hours of semi-structured interviews with 8 users. Interviews ranged from 30 minutes to 3 hours. We also analyzed four blogs and two web forums associated with the formal activist community.

INFRASTRUCTURES OF LOOKING IN SECOND LIFE

Past research has often focused on avatar embodiments and virtual objects in a digital representation of three-dimensional (3D) space in their analyses of CVE. Supplementing this understanding, our observations indicate that SL has developed a much more complex world of alternate information displays, profile pages for avatars, and application programmer interfaces (APIs) that allow end users to create custom objects that can sense and act.

¹ Genderqueer is an identity category that is inclusive of all who do not identify with traditional male and female gender identities.

Understanding these complexities and the practices that develop around them in real world CVEs can deepen our understanding of the potentials and design challenges involved in creating these environments.

Understanding practices of looking in SL requires a detailed understanding of the user interfaces through which people connect to and navigate in SL, including the interdependencies between apparent gaze and the world, as it is visible to the end user. Affordances of looking – both the capabilities of looking and the visibility of those practices to others – depend not only on the design and the user’s interpretation of the CVE, but also the idiosyncratic, unpredictable quirks and bugs that emerge when designs meet unexpected deployment conditions. These affordances become resources and constraints for SL residents’ awareness of their environment – an awareness grounded in learned practices of seeing and specific to local needs and concerns. This awareness is the foundation of managing one’s audience, shaping one’s presentation of self in everyday Second Life. In this section, we consider the implementation and technical possibilities of looking, being seen to look, and being the subject of “looking” in SL.

Layered Infrastructures of Looking

Second Life defaults to third-person point of view. This perspectival strategy provides users with views of the world with their avatars in it, as well as a wider field of vision than allowed by first-person point of view [15]. Each user sees through an implied camera placed at some radius from the avatar body. By zooming in and out, the radius of the camera around the avatar increases, providing broader or narrower vistas onto the virtual world. Using the arrow keys, users can move the camera along a sphere defined by the sweeping radial camera distance from the avatar, allowing them to view the world around them without reorienting their avatar bodies.

A first-person perspective, called “mouselook,” provides a view on the world as though through the eyes of the avatar, such that the avatar body is no longer visible. This view is analogous to the physical world perspective of unmediated looking. In this perspective, users must actually turn and reorient themselves to see the environment behind or to the sides of them. The major difference in capabilities between this view and the “real world” physical viewing capabilities of human beings is mouselook’s lack of peripheral vision as compared to “real world” vision.

Beyond the potential for looking enabled by the camera, the SL screen, like those of many virtual worlds, layers additional visual information in a heads up display (HUD) (See Figure 1). A miniature map shows a topological overview of one’s currently occupied region with green dots representing locations of other residents. Person-to-person instant messaging (IM) and chat text from nearby residents visually appear overlaid on the user’s view of the world. Because the visual interface in Second Life includes

“hearing” utterances, people often mix vocabulary of seeing and hearing or looking and listening in SL. For example, one informant described speaking privately in SL: “If I’m *discussing* [something] very personal, I always do so in IM so that no-one else can *see*” (emphasis ours). Thus, in our analysis of looking in SL, we consider verbal, written expressions, as well as abstracted and supplementary information layers, alongside the more commonly considered views of the 3D world.




Figure 1: Third-person perspective, overlaid with camera controls and a mini-map. Menus are along the top and bottom of the screen.

Looking Beyond the Simulation

Taking this notion of layered infrastructure a little further, looking in virtual worlds extends beyond the 3D simulation. While discussions of interaction in a virtual environment such as SL inevitably focus on the simulated world as both the site and the form of interaction, SL is simultaneously a graphical user interface and potentially just one of many active applications. Thus, interaction with, through, and around that application interface is also relevant to any account of looking and examining the world.

Beyond the simple forms of virtual looking described in the previous section, the most common and easily executed act of looking is to examine another avatar’s “profile” – a list of properties and descriptions set by each user. The simplest way to view the profile of a copresent avatar is to right-click the avatar and select “Profile...” from the resulting context menu. However, there is an interesting dualism at work here, in the way that the mouse is being used both as a pointing tool in the virtual world *and* as a pointing tool within the graphical interface. Because avatar gaze in SL coarsely follows the mouse pointer, the action of activating an avatar profile necessarily redirects the looker’s gaze towards the examined other. This particular mode of interaction can signal curiosity and may be taken as an opening for conversation. In contrast, users can surreptitiously investigate a copresent Second Life resident’s profile by using “Search” to find the resident by name. Keyboard actions are less accountably visible within

the user interface. Avatars of users looking at profiles merely appear idle. It is therefore some of the means of obtaining the profile, rather than the looking itself, that are rendered visible, thus making some profile viewings accountable passively or automatically.

Another example of accountable looking is taking a screen capture of the virtual world scene using the “Take Snapshot” command. When this command is executed, the avatar raises an iconic camera to its face and a shutter makes a sound. Reinforcing visual and audio signals make the moment of capture explicit. To those who may be captured or “photographed” in this manner, the behavior of the avatar provides a lightweight account of the recording as a resource for understanding and negotiating these activities.

A third way to look is to utilize the “Alt”+click tool for looking at specific objects. “Alt” + click effectively trains the looking avatar’s gaze at the “looked-at,” centering the view on that individual and directing the looking avatar’s head to the looked-at. However, moving the mouse disengages the gazing, directed head, moving the head again to follow the mouse.

Although these means of looking are accountable, they are often only discernible with difficulty, even for experienced participants. After someone had just taught the maneuver to one of us, we tried it a few minutes later to see if he noticed:

Me: “I just tried the alt+click to look at you”
 Me: “did you notice? :-D”
 Dwayne²: “:) probably not. at least consciously”

Although there are too little data to understand completely the effectiveness of accountable gazing, some challenges are already clear. The mouse moves the avatar head, redirecting the apparent gaze, but the mouse may move for many reasons having nothing to do with looking. For example, the user may be mousing between Second Life and another application, such as a web browser. The user may be using the commands at the bottom of the screen. Noisy head movements threaten to render even the most explicit, intentional gaze unintelligible.

What is particularly of interest here, though, is the ways in which these different modalities of looking – those that draw upon the conventions of the virtual world, those that draw upon the conventions of the graphical user interface, and those that draw upon the conventions of a search engine – are intertwined. They cannot easily be separated. Interaction in the virtual world extends beyond the confines of the simulation window. The virtual world manifests itself

² All participant names are pseudonyms.

here as one part of a user interface that incorporates multiple modes of interaction.

LOOKING AS VISUAL PRACTICE

In this section, we examine the practices of looking in more detail. These practices rely upon a range of technological features, of course, but here focus on how those technological features are deployed. Out of these features, interwoven with the specificities of other virtual world infrastructures, users develop local practices to achieve their embodied, social needs.

Built-in Looking

First, we consider the many ways users can intentionally visually explore the SL virtual world. As noted in the description of the infrastructures of looking, the (invisible) camera that sees the virtual world can be moved along the (invisible) sphere that surrounds the user's avatar body. Using advanced camera controls, the camera can also be moved to arbitrary but finite points in the 3-D space around the avatar. This tool for looking allows the user to look where his or her CVE body is not. For example, the camera can be moved through walls, floors, and ceilings. As a notecard of security advice written by one SL user explains, "clever manipulation of the camera angle (by anyone) ... can sometimes expose more than you thought to explorer's view." The movement of the camera in space, combined with zooming, which narrows or expands the viewable area, gives the user great flexibility in focusing their view of the world. Yet other users cannot detect or surmise another user's viewable area. Although these problems of intersubjective understanding, as well as their possible design solutions, have been identified by CSCW researchers [15, 16], SL users cope with this mode of unaccountable looking in their everyday lives.

One of the first things that most users do in SL is to "deruth" oneself, that is, to change one's avatar from the default body to a personally-chosen appearance. Recent SL clients provide a specialized view that supports this, but in earlier versions, this activity required that users learn how to reposition the camera so that they could see their own avatar more clearly. While giving clothing to one of us, an experienced SL informant indicated a lengthy menu of advanced camera controls in response to the recipient's complaints about assessing her avatar's appearance from different angles. These camera controls become an important resource in social interaction by supporting self-perception, including the ability to visually perceive oneself as others do.

The movement of the camera can also be used to view all participants in a conversation while still maintaining socially acceptable proxemics. In two situations, we observed the camera being used in just this way. In both cases, informants sat in a meeting room where chairs were arranged in a circle, drawing from physical world models of discussion seating arrangements. The limited gestural cues

and narrow default field of view, however, meant that the symbolically meaningful seating arrangement, signaling cultural values of participation and non-hierarchical organization, suppressed important cues of social ordering in the virtual world. To restore their awareness of social ordering cues, such as moving hands to indicate typing, informants described moving the camera to overhead views of the room or other angles that let them view those who were present.

Cameras have also replaced body movement in some situations in which physical body movement would have been necessary in the "real world." In one case, we discovered a SL resident hidden behind the walls of her virtual home peering around the outside of the home by moving her camera through the building's walls. This behavior became apparent when a disembodied voice greeted one of us by name. The peeking resident explained that she was looking outside to see if she had left a certain object behind. Instead of walking or flying her avatar outside of her home, she chose to peer remotely instead. Several end-user tool developers in SL have taken this visual mobility a step further by creating "camera follower" scripts that allow users to navigate their cameras easily to their desired destination and then teleport, or instantly transport their bodies to that point. One developer advertises the tool as a way to get into places with no doors and to "outrun followers" – "walls, floors, ceilings, nothing is in the way." These practices illustrate how the drive for convenience or expediency can motivate creative end-user appropriations of the technological environment. This allows for the development of behaviors and uses not envisioned by Linden Lab's designers and engineers. In this case, the challenges of moving an avatar around obstacles and recovering from collisions are overcome by travel to a destination using a mobile and flexible gaze.

End-User Augmented Looking

The world's scripting language, LSL [20], provides powerful mechanisms for creating tools to augment looking. Among the many objects created and given away or sold are tools that mediate and extend SL residents' layered vision. In this section, we examine these tools for what they do in action, symbolically, interactionally, and otherwise, and what that tells us about practices of looking.

Experienced informants – six in total – acquired tools, such as *Runemaster* and *Mystitool*, to augment their default capabilities of seeing. Bobbi, who plays as a child in SL, described special glasses she purchased both as ornament and perceptual augmentation in the virtual world: "My glasses come with a sensor so I know who's around...within 90 metres [in Second Life] of where I am." She also has a package of scripts called *Runemaster* that includes among its functions the ability to list the names of all Second Life residents in the currently occupied region³.

³ A region is a named unit of land in Second Life.

Four other informants described using *Mystitool*, an end-user created kit with similar functionality to *Runemaster*. One informant also described creating his own, custom tool that would indicate who is in “chat distance.”

Rather than providing a persistent display, both tools require the user to issue a command to get a list of people nearby, either within shout range or in the same region. Some tools list the names and distances of nearby avatars. Four informants described constantly reading those distances against their knowledge of the distance chat and shouting travels – 20 and 96 (virtual) meters, respectively. In addition to the common practice of detecting which avatars were nearby, people also reported using these tools to verify who might hear before talking and to find other people in a crowded area.

Some users of these tools described pervasive needs for such tools of augmented seeing. A web forum poster wrote: “I cannot imagine not having it – I use it constantly and I use it every day.” Similarly, another informant reported looking at the tool just seconds prior to explaining, “I look at it pretty much constantly. It’s just part of my world awareness.” This checking requires executing text commands, rather than the sorts of lightweight awareness advocated by the CSCW community, but the need for awareness appears to be compelling enough to some users to justify the interactional load.

It is telling that these artifacts are *designed* to be unaccountable. Several users of *Mystitool* chose words such as “unobstrusive” and “tiny” to praise the tool on an online forum. Yet these devices were also openly discussed. This, and the prevalence of these sensing devices among more experienced informants, suggests that unaccountable mode of looking is pervasive and out of the way, rather than intentionally surreptitious.

INFRASTRUCTURE AND IMPLEMENTATION

In this paper, we demonstrate that the varied practices and mechanisms of looking in SL cannot be separated from the infrastructure and implementation context within which they exist. Looking, then, is not merely a property of a virtual world or simulated reality; it is a property of a particular system that combines different interactional modalities, offers specific technological features, and runs on particular kinds of computers using specific forms of interface.

Furthermore, we were struck by the importance not just of *design* “beyond the virtual,” but by the relevance of specific artifacts of the *implementation*. That is to say, the technological choices made by designers, or the constraints afforded by particular network infrastructures and hardware capabilities, are not aspects of SL’s interaction design. However, they most certainly are aspects of how SL appears to users – the user *interaction experience*.

Through breakdown, hardware and network infrastructures reveal themselves as consequential in looking. These

infrastructures impact the graphical world’s speed of rendering, effectively determining short- or near-sightedness in SL. For example, more distant objects render late and rendering a large radius in real-time slows performance. Thus, some users set SL’s rendering limits to a distance less than the default, limiting their range of vision.

The most unpredictable instances of unaccountable, and even unintentional, looking come from these particularities in the graphical rendering process. These quirks result in perceptual asymmetries in which occupants of the same virtual environments have no way of telling what others in their social space can see. The problems of such perceptual asymmetries are well-documented in CSCW literature [4, 14, 15].

As a particular example of these quirks, slow rendering speed can provide fleeting views inside walls and under clothes. When an avatar enters a not-yet-rendered area, the user sees the region rendered progressively, plane-by-plane. During an observation session in this research, a house with two occupants lying on a bed progressively rendered with the back wall first, then bed and occupants, then sidewalls, and only then the obscuring front wall and roof. The occupants of the house could not have recognized their vulnerability to scrutinizing eyes, because they were not moving, and thus their rendered surrounds appeared stable. In several other instances, avatar appearance stopped refreshing in the middle of an outfit change, leaving participants perceiving themselves as topless or with parts of two disparate costumes. In these cases, the only remedy was to find a trusted person and query about one’s own appearance.

These unreliabilities of rendering destabilize the foundations of self-perception that ground social behavior. The asymmetry between the perception of self and other undermines a principal promise of virtual worlds in collaborative social environments – that, according to Hindmarsh *et al.*, problems of establishing what another can see are overcome through sharing “views of a stable and common virtual world” [15].

As another example, technological quirks of the SL client undid the expectation that a label is fixed to that which it identifies. One might naturally expect a label to be part of, and thus inseparable from, the visual representation of an avatar. In fact, it is actually an *extension* of that visual representation, and as such, may be present in areas in which the virtual “body” is not. For example, during one of our explorations, one of us came across a large floating spaceship. The doors were locked, but this quirk of Second Life’s rendering engine showed the identity labels of two ship inhabitants floating outside. The labels appeared above the ship because it was being viewed from far enough away that the bubbles had to grow to be legible (See Figure 2). The occupants may have also been near the ceiling. This is a sufficiently uncommon experience that it would be

difficult and perhaps pointless to check actively whether one is running the risk of having their presence projected beyond the walls that enclose him or her, but it is illustrative of the potential visibility of action brought on by invisible forces, such as a rendering speed and unexpected effects of information design decisions.



Figure 2: The names of occupants inside the space ship are rendered outside, unbeknownst to the occupants.

Here, we have seen how common and unavoidable ways of looking in SL are unaccountable to social actors occupying the virtual world. Because the camera is unintelligible, social actors have no way of discerning what an avatar's viewing range might be. Custom-built objects empower actors to surreptitiously monitor the presence and location of others in a region. And quirks of graphical rendering fleetingly but visibly expose sites occupants believed to be hidden.

More importantly, however, looking is by default unaccountable and thus can create difficulties in social interaction. For example, it is difficult to signal to others that one is *not* looking in situations where politeness or other social mores may necessitate this signal. Being looked at by those hidden behind walls or far away engenders the sense of being in Bentham's Panopticon, where one cannot know when they are being looked at – only that they always may subject to the gaze. Spaces of varying, accountable levels of exposure are a resource for building intimacy and bonds, but, conversely, spaces that are all similarly exposed impoverishes actors of expressive resources, such as secret hideouts “for your eyes only” or the public places one goes “to see and be seen.”

MANAGING AUDIENCE

Ultimately, the issues surrounding looking in this visual interface to a virtual world shape practices of self-presentation and audience management. Managing audience consists of maintaining awareness of others' presence and choosing a performance space relative to the understood and desired audience. Performance spaces may be in-world

(e.g., through IM or SL locations) or out-of-world (e.g., web forums or face-to-face).

Boundary Management On the Fly

Boundary management can include *ad hoc* and highly situated negotiations of what “to share and not to share” in the flow of other activities [1]. Tools for “world awareness,” as one informant described them, are critical infrastructure used to perceive, monitor, and adjust to changing audiences in the virtual world.

Anna, an autism activist who designs and implements group spaces for various disability support groups, described the challenges of perceiving context and negotiating boundaries in the infrastructure of SL. Designing a space for an anorexia support group, Anna built a virtual coffee table that displayed the names and distances of nearby avatars (see Figure 3). This approach attempted to support presence awareness as a way of letting participants shape their communications with respect to the audience in the vicinity.




Figure 3: End-user built coffee table with presence awareness.

She and other organizers, however, ultimately decided the table failed to foster the level of safety that group participants required to speak freely. Anna located the problem in the frequency with which the scanner detected nearby avatars: the presence monitoring code, called an avatar scanner, was only able to detect presence every second. Thus, it was possible for a conversant to type a lengthy message assuming the presence of certain avatars, only to send the message just after a new person entered the conversation space as of yet undetected by the scanner. The message will be visible to those present at the time of transmission, not composition. While Anna felt that this table was in excess of what her autism activist group needed, it was insufficient for the anorexia support group.

Although audience management is pervasive, it is just one of many concerns that are juggled, foregrounded, and backgrounded in the course of an interaction. Over the course of an interview that began about security issues in SL, the conversation moved to topics such as Dwayne's history, relationships, and space in SL.

Dwayne: "Is this conversation under interview rules per the human-subjects form, or are we just chatting? :)"

Me: "it can not be part of the research"

Dwayne: "Just want to make sure of the shield of confidentiality."

Me: "Oh yes, definitely."

Dwayne then began to describe his history of romantic encounters in Second Life, as well as his experiences of sexuality on the wider Internet. Sitting outside of his Second Life home, other people's distant shouts occasionally interrupted us, as we chatted over the "local chat:"

Me: "oh, I just realized we're not IMing."

Dwayne: ":) that's true."

Dwayne: "No one in chatrange, but there could be listening devices. :)"

Dwayne: [after switching to IM]
"Good idea. :)"

Newly conscious of our surroundings, we switched to private in-world IM to continue the conversation. This example foregrounds the challenges and everyday difficulties of managing audiences in the flow of everyday activity. Dwayne is in some ways an expert informant – he understands and uses avatar scanner tools and he described the possibility of invisible listening objects. Yet, even in his account, we see that in the flow of interaction, especially when private and public channels of communication look identical in the interface, maneuvering between these modes reliably can be challenging. Furthermore, the possibility of unaccountable looking through listening devices makes Dwayne's true audience unknowable. Both our and Dwayne's technical sophistication is insufficient to guard the confidentiality he carefully sought. We, together, assessed the situation and came to a mutual decision to switch communications channels. Boundary management here is collaborative, ambiguous, imperfect, and grounded in the practices of looking in the virtual world.

Disappearing Acts

Hiding, sometimes with others, is one form of negotiating boundaries and managing audience. In SL, the many unaccountable modes of looking mean almost anyone may be anywhere at any time. These strategies rely on infrastructures of space and embodiment, as well as understandings of the what those being avoided can and cannot see – understandings often grounded in one's own experience of the world.

Dwayne, an informant who prefers to experiment with his avatar's appearance away from "prying eyes," described

seeking "probable temporary privacy" by flying high up into the empty sections of sky where passersby are unlikely to wander or move their cameras. An advanced user, Dwayne qualified his hiding technique, explaining that avatars can fly to great heights when equipped with "flight assist" scripts, and they can also use tools to detect his avatar in a region. In seeking "probable privacy," Dwayne came to a workable if not ideal way of hiding as a solution based on his resources and assumptions about others' desire or incentive to interrupt him.

Similarly, after her presence awareness coffee table proved insufficient, Anna came to the conclusion that hiding would be the only way to achieve sufficient privacy for the anorexia support group participants. Her conclusion, established from her own experiences and those of group organizers, is similar to that of medical researchers who concluded that skyboxes with full security systems would be appropriate solutions for medically sensitive groups [5]. A skybox is a building erected high in the sky, isolated from the casual passerby. Skyboxes also repel any unauthorized persons who come within chat distance to another SL region. This approach and the effort required to implement it demonstrates the importance of hiding in certain situations.

DISCUSSION

Working from the results of this work, in this section, we examine fundamental themes our observations capture and the broader relevance of these considerations. In particular, these data reveal the variety of ways of looking and their integration into a range of visual and monitoring practices within the virtual world. Some of these mimic the conventions of everyday spatiality; others are firmly situated within the realm of conventional desktop interfaces; some are consequences of infrastructure and implementation, and so are uniquely tied to the particulars of SL; and still others are hybrid practices conditioned by all of these considerations at once. Although it is tempting to construe "looking" as merely the accountable or observable gaze of the avatar, we chose to examine this phenomenon in SL more broadly.

Over the last fifteen years, CSCW researchers have critiqued mediated interaction in various forms, including in CVEs, for the communication asymmetries they introduce and their failures to replicate those elements of everyday interaction that support smooth coordination of activity [14, 15]. Certainly, our data supports many of these points. Furthermore, the forms of communicative breakdown and compensatory behaviors we observed point to the failures of conventional 3D simulations to support everyday interaction.

Alongside these critiques, however, one must also consider three new observations that emerged from this study. First, these environments provide people with resources for reconfiguration and repurposing that go beyond what everyday interactional environments may offer. Second, the

“simulation of reality” plays only a partial role in coordination, collaboration, and interaction in CVEs, which need to be seen simultaneously as graphical and textual interfaces. Third, new patterns of engagement arise over time and create a context for interaction that is uniquely tailored to the implementational specificities of the system in question, rather than simply attempting to recapture the effective properties of the everyday medium.

Perhaps most interesting is the fundamental challenge of making broad statements about interactional practice in environments that are as rich and varied as SL. Indeed, we take this to be the primary implication of any examination of situated practices of looking – they are situated not only within particular systems or implementations, but also in particular bodies of practice that arise around them.

Community Practices

We draw considerable inspiration from Goodwin’s examination of “professional vision” [13]. Goodwin highlights how various forms of scientific and professional practice are intimately connected to ways of seeing the world and organizing it into socially accountable objects and events. The visual availability of these objects and events is constituted through forms of practice shared by professional groups. For example, Goodwin describes the ways that jurors are taught to view the Rodney King video like a police officer, learning to discern what poses that an officer might read as threatening poses but a civilian might not interpret the same way. Goodwin explains how archaeologists similarly develop professional practices of seeing that undergird collaboration. It is critical, however, that the phenomenon of interest made discernable by each community’s practices are unique. In other words, professionals work towards mutual intelligibility by pointing to phenomena of interest, but those phenomena vary and are made discernable through unique practices. In an altogether different domain, our results also demonstrate how a body of shared practice renders forms of looking and acting visible to others as a basis for interaction. In Goodwin’s analysis, professional vision may be mediated by a range of representational and technological systems; so here too. However, whereas in Goodwin’s case, the technological systems are relatively stable and the professional practices well established; in our examination, the technology and its consequences are still fluctuating and developing on both a broad scale as the technology itself changes and a local scale as individual user practices develop around it.

In this paper, we reveal some of the ways users are coming to grips with these phenomena and understanding their consequences. However, broad statements about “interaction in Second Life” are unlikely to be tenable. Rather, our investigation indicates that any future work, analytic or design-focused, must support the kinds of diverse localisms that characterize phenomena in continual flux and development.

For example, this work focused on a set of users whose reactions to and considerations of privacy and visibility varied across dimensions such as permanence and isolation. In many ways, the importance of these dimensions for these communities turns on their identification and activities as focused on minority or disabilities related issues. In other communities, the needs of the community members would likely turn on other dimensions, for example symbolic safety or level of technical expertise. Thus, rather than attempting to codify these dimensions, we offer a consideration of looking, visibility, and privacy as situated sensitizing mechanisms for designers and empiricists alike.

Individual Evolution

In addition to practices being highly localized and situated in terms of community, they are also temporally local within even a single individual. Less experienced users cannot see through the eyes of the more experienced. Experienced users remember being new, but less and less as time goes on. Thus, like an expatriate in a foreign land, people over time adapt to and appropriate the customs of this space - whether these customs be ignoring sometimes stuttering gaze or comprehending the oddities of a not yet fully rendered space. The familiar becomes strange and the strange familiar as the world itself changes, they change within it, and perhaps their hardware and surrounding infrastructure changes as well.

CONCLUSIONS

Despite arguments and intuitions that interaction in virtual worlds may be smoother, more “natural” or more “intuitive” than interaction in other mediated contexts, several authors in CSCW have pointed to the problems inherent to these kinds of simulations [15,22]. These problems include the failure of virtual worlds to support various aspects of interaction. In this paper, we also argue to be cautious of the line of reasoning that implies interaction in virtual worlds is superior due to its simulation of everyday reality but for a different reason. The interactional site in virtual worlds is not simply a simulation of everyday reality. Rather, it is simultaneously a graphical and textual interface, one that is further inflected by the implementational realities of the platform on which it is encountered. While virtual worlds may seem to offer, by analogy, certain benefits of “real world” interaction, virtual worlds, details of virtual embodied experience offer distinct challenges and opportunities.

We should note that we do not think of this, or present it, as a problem with virtual worlds *per se*. Rather, the problem lies only in certain interpretations of what virtual worlds can do and why. Indeed, in this work, we describe a range of creative ways the system can be and has been appropriated to create a novel platform for interaction. This platform plays with certain metaphors of everyday interaction but is by no means constrained by them. Allowing for this sort of appropriation seems both positive and fitting. At the same time, it suggests that we as a

community should further be skeptical of broad claims about the interactional properties of virtual worlds, because the meaning of even such seemingly obvious acts as looking and seeing is, itself, embodied in highly particular systems and their emergent practices. Thus, these topics – emergence and appropriation – are perhaps even more useful topics for examination and targets for design than generic mechanisms for interactional support.

ACKNOWLEDGMENTS

We'd like to thank Silvia Lindtner and our anonymous reviewers for their helpful comments. Thanks to Rosemarie Garland-Thomson for partly inspiring this work with her article "Ways of Staring." This work was supported in part by the National Science Foundation under awards 0527729, 0524033, and 0712890.

REFERENCES

1. Anderson, K. and Dourish, P. Situated Privacies: Do you know where your mother [trucker] is. *Proc International Conference on Human-Computer Interaction*. (2005), 22-27.
2. Antonijevic, S., From text to gesture online: A microethnographic analysis of nonverbal communication in the Second Life virtual environment, *Information, Communication & Society* 11, 2(2008), 221-238.
3. Bailenson, J., and J. Blascovich, A.C. Beall, and J.M. Loomis, "Equilibrium theory revisited: Mutual gaze and personal space in virtual environments," *Presence: Teleoperators and Virtual Environments* 10, 6 (2001), 583-598.
4. Benford, S. et al., User embodiment in collaborative virtual environments, *Proc. CHI 1995*, ACM Press (1995), 242-249.
5. Boulos, M.N.K., Hetherington L., and Wheeler, S., Second Life: an overview of the potential of 3-D virtual worlds in medical and health education, *Health Information and Libraries Journal* 24, 4 (2007), 233-245.
6. Brown, B. and Bell, M., CSCW at play: 'There' as a collaborative virtual environment, *Proc. CSCW 2004*, ACM Press (2004), 350-359.
7. Dourish, P., Adler, A., Bellotti, V. and Henderson, A., Your Place or Mine? Learning from Long-Term Use of Audio-Video Communication, in *Computer-Supported Cooperative Work* 5, 1 (1996), 33-62.
8. Dourish, P. and Anderson, K., Collective Information Practice: Exploring Privacy and Security as Social and Cultural Phenomena, *Human-Computer Interaction* 21, (2006), 319-342.
9. Garau, M., Slater M., Bee, S., Sasse, M.A., The impact of eye gaze on communication using humanoid avatars, *Proc. CHI 2001*, ACM Press (2001), 309-316.
10. Garfinkel, H. *Studies in Ethnomethodology*. Englewood Cliffs, NJ: Prentice Hall, 1967.
11. Gaver, W.W., The affordances of media spaces for collaboration, *Proc. CSCW 1992*, ACM Press (1992), 17-24.
12. Goggin, G. and Newell, C., *Digital Disability: The Social Construction of Disability in New Media*, Rowman & Littlefield Publishers, Inc., 2002.
13. Goodwin, C., Professional Vision, *American Anthropologist* 96, 3 (1994), 606-633.
14. Heath C. and Luff, P., Media Space and Communicative Asymmetries: Preliminary Observations of Video-Mediated Interaction, *Human-Computer Interaction* 7, 3(1992), 315-346.
15. Hindmarsh, J., Fraser, M., Heath, C., Benford, S., Greenhalgh, C., Fragmented interaction: establishing mutual orientation in virtual environments, *Proc. CSCW 1998*, ACM Press (1998), 217-226.
16. Hindmarsh, J., Fraser, M., Heath, C., and Benford, S., Object-focused interaction in collaborative virtual environments, *ACM Transactions on Computer-Human Interaction* 7, 4 (2000), 477-509.
17. Hindmarsh, J. and Heath, C., and M. Fraser, (Im)materiality, virtual reality and interaction: grounding the 'virtual' in studies of technology in action, *The Sociological Review* 54, 4 (2007), 795-817.
18. Hine, C., *Virtual Ethnography*, Sage Publications Ltd, 2000.
19. Hollan, J. and Stornetta, S., Beyond being there. *Proc. CHI 1992*, ACM Press (1992), 119-125.
20. LSL Portal - Second Life Wiki. http://wiki.secondlife.com/wiki/LSL_Portal.
21. Marcus, G.E. Ethnography in/of the World System: The Emergence of Multi-Sited Ethnography, *Annual Review of Anthropology* 24, (1995), 95-117.
22. Moore, R.J., Ducheneaut, N., and Nickell, E., Doing Virtually Nothing: Awareness and Accountability in Massively Multiplayer Online Worlds, *Computer Supported Cooperative Work* 16, 3 (2007), 265-305.
23. Taylor, T.L. Living Digitally: Embodiment in Virtual Worlds, *The Social Life of Avatars: Presence and Interaction in Shared Virtual Environments*, 2002.
24. Yee, N. and Bailenson, J., The Proteus Effect: The Effect of Transformed Self-Representation on Behavior, *Human Communication Research* 33, 3(2007), 271-290.