

Spatial Social Behavior in Second Life

Doron Friedman^{1,2}, Anthony Steed¹, and Mel Slater^{1,3}

¹ University College London,

{d.friedman,a.steed,m.slater}@cs.ucl.ac.uk

² The Interdisciplinary Center, Herzliya, Israel

³ ICREA-Universitat Politècnica de Catalunya, Spain

Abstract. We have developed software bots that inhabit the popular online social environment SecondLife (SL). Our bots can wander around, collect data, engage in simple interactions, and carry out simple automated experiments. In this paper we use our bots to study spatial social behavior. We found an indication that SL users display distinct spatial behavior when interacting with other users. In addition, in an automated experiment carried out by our bot, we found that users, when their avatars were approached by our bot, tended to respond by moving their avatar, further indicating the significance of proxemics in SL.

1 Introduction

Social virtual environments, as envisioned in cyberpunk literature [12, 4], are now becoming widely popular. At the time of writing these lines the SecondLife (SL) web site brags there are approximately 5.9 million registered users and 25,000 online visitors⁴. SL is not the first such online 3D persistent community, and it is by no means flawless; we are definitely going to see SL and rival products change and evolve significantly. However, SL's increasing popularity have already made it an interesting target of research in its own right; it is an interesting opportunity to study the behavioral patterns of people in such virtual universes.

In highly immersive virtual-reality (VR) environments we have evidence that participants have a strong sense of presence; at least some of them for some of the time. We follow Sanchez-Vives and Slater's operative definition of presence [17] in that we expect people to behave in the same way in the virtual environment as they would in an equivalent real-world situation. Note that presence is, thus, very different from engagement. There is no question whether SL is engaging, our question is: do people have a sense of presence in SL? do they behave as if they were in a real world, and to what extent?

One of the main features distinguishing SL from online chat environments is that it is supposed to induce a sense of being in a three-dimensional space. Therefore, we set out to study the spatial behavior of SL users, and specifically, their social spatial behavior.

Since SL is a commercial product, we do not have access to statistical data, and it is unlikely that such data would be made public by Linden Labs (the

⁴ <http://www.secondlife.com>

company behind SL). Thus, we have developed automated software bots that are able to wander around SL and systematically collect data. In addition, we have built these bots to have some social capabilities of their own; this allows them not only to observe and collect data, but also to participate in social interactions, thus essentially carrying out social experiments within SL.

2 Background

The study of the social significance of space, and in particular the distance between people as they interact, was initiated by Hall [5, 6], who termed it proxemics. Hayduk [7] provides a more recent survey.

Video games have been recognized in the past as a potential for artificial intelligence research (for example, see [10] and [9]). As Loomis et al. [11] note, collaborative virtual environments (CVEs) are also a useful tool for research in psychology.

CVEs provide each user with an avatar that embodies them in the virtual world. This avatar serves several purposes: it should be distinctive to convey identity, at least over a short period of time, it indicates the user's chosen position, and possibly their attention and some indication of their emotional state, through verbal or non-verbal communication (NVC) [2].

However, almost exclusively in these environments, the user must "act" at least some of these: they must chose their appearance, position their own avatar and indicate their emotional state. Depending on the system, some of this may be automatic or semi-automatic. In an immersive system, or other tracked system, some NVC may occur automatically, because some of the participant's limbs are tracked. In a desktop system, attention and position are conveyed semi-automatically because the nature of the interface means that the user must normally get their view close to an object to interact with it. Some CVEs support 1st-person views, but most support 3rd-person views, or 3rd-person is the most prevalent view if there is an option. In fact, some CVEs have very rigid constraints between the 3D view and the avatar position. SL is slightly looser in this respect, in that the user can rotate and zoom the camera around their avatar's location.

The avatar thus "grounds" the user in the environment. A CVE system may enable communication between users if their avatars are nearby, though this is typically determined only by proximity, not by facing direction. Despite this, experience with early CVEs found that users would naturally form social groups and face users they were communicating with [3]. Even experienced CVE users, who know that avatar proximity is neither necessary nor sufficient to enable communication, adopt socially-aware spatial behaviors: these behaviors help manage what could otherwise be quite confusing situations with many people attempting to interact at once [13]. Several studies of immersive systems have found evidence that users treat social space very naturally, attending to the gaze of others and watching body language [8].

Recently, Yee, Bailenson and colleagues have studied social space within SL. In [18] they uncover patterns of social space use that would be expected, such as gender differences, and eye-gaze avoidance for situations where the interpersonal difference is only 2-4m.

The previously discussed studies were all conducted with avatars of users. Bailenson et al. [1] studied approach of an immersed user to both agent avatar and user avatars. They found that users would keep distance, but that knowledge that the avatar was a user avatar would deter personal space invasions. Vinayagamoorthy et al. [15] studied the situation of a user of an immersive system approaching agent avatars, which were programmed to represent different emotional states. The agent avatars would respond when the user approached, and the user would subsequently adopt a socially-acceptable position. Comments from interviews, indicated that some users felt they should respect the social conventions, even though they knew the avatars were autonomous [16].

3 The SL Bots

SL is intended to be built first and foremost by its users, and it thus provides facilities for content creation. Programming is achieved with the Linden Scripting Language (LSL), which provides a wide range of capabilities; at this time LSL includes 330 built-in functions, including: vehicles, collision detection, physics simulations, communication among users, inventory management, playing audio and video files, and more. However, LSL was clearly not designed to construct bots; scripts are only attached to objects, not to the user's avatar directly. We have come around this limitation by attaching a ring to our avatar. The ring object can then run a script, and the script can then be used to move the avatar and animate it, so that it appears walking while moving, as well as performing other tasks required by our bot.

Our bot has a basic capability for wandering around and finding locations or objects of interest. The implementation is as follows: the bot selects a random direction and starts walking in that direction, until it either reaches an obstacle (such as a wall) or the target. If it reaches an obstacle it selects another random direction and keeps moving in the new direction. While this approach is simple and not necessarily efficient, it has proved successful in practice, and even allowed our bot to occasionally wander in and out of closed buildings, passing through doors. Typically, as in the study described here, the bot is instructed to locate other avatars. When it detects one or more avatars it stops and carries out its social task, until it is terminated, or until it find itself alone again.

The bot has simple interaction capabilities: when it encounters other avatars it greets them using their name. The bot can also play a large range of approximately 50 pre-recorded animations. However, we have found out that such animations do not play an important role in SL. While users are able to allow their avatar to play pre-recorded animations, this is not similar to real-life NVC. For example, users very rarely use these animations, and when they do, these are typically high-level animations, such as dancing. In the real world, NVC is

a continuous process, which plays an important part in communication, and is mostly unconscious.

Our bot has capabilities for data collection: it can be instructed to collect information about the objects and avatars it encounters on the way, log this information, and send it to us. Currently it uses email to send us the information, but SL allows other forms of communication with external software, such as HTTP or XML-RPC⁵.

4 Experiment 1: Proximity in Dyadic Interactions

SL adopts the approach, typical of non-immersive VEs, of conversational characters [14]: the avatars display autonomous NVC, with a possibility for the users to override their avatars' gestures and postures. In SL, users very rarely use this possibility. Such autonomous NVC may or may not contribute to the sense of presence of SL users, but it clearly does not allow us to study the level of presence based on the users' behavior. Thus, instead of NVC, we study a subset of NVC, namely proxemics.

Most proxemics research is focused on dyadic interactions (interaction between two people). Hall [6] distinguished among several distance categories, measured as circles around a person:

1. intimate space: for touching — up to 1.5 feet.
2. personal space: for interaction among friends — 1.5-4 feet.
3. social space: interactions among acquaintances — 4-12 feet.
4. public space: for public speaking — over 12 feet.

If indeed users copy this social behavior into SL, we expect to see similar distances among interacting avatars. SL avatars are, by default, of the same height as average adults (measured, of course, in virtual units). While users can create very small or very large avatars, it is rare to see such cases⁶. Thus, we expect the virtual distance in SL to match the corresponding real distance in a social real-world interaction, with real units replaced by virtual units. Generally, SL does not allow touching (with the exception of unique devices for “adult” interaction, which were not studied here). We thus do not expect to see avatars within intimate space.

4.1 Method

We have sent our bot on a mission to collect spatial data in SL, from 20 different locations, selected arbitrarily. It is not possible to follow Yee et al. [18], who sampled all SL regions, because: i) SL is growing fast, and ii) some areas in SL

⁵ <http://rpgstats.com/wiki/index.php?title=XMLRPC>

⁶ The same goes for non-human avatars; while it is common to see humanoid avatars with fancy or unique clothing, it is quite rare to encounter non-humanoid avatars (such as animals).

block our script from running. We are still looking for a method to ensure that our sample of SL is balanced. Similar to Yee et al. [18], we have excluded from the study specific regions such as dance clubs, sex clubs, or other locations that may pose special constraints on spatial behavior.

From this data, we first isolated cases of dyadic interactions. Two avatars are considered interacting if they were in the same area for over one minute and if they were facing each other. By facing each other, we mean that their orientation was no more than 90 degrees away from the line connecting their positions.

Note that, unlike Yee et al. [18], we only look at couples of avatars that are alone (i.e., all other avatars are over 10 meters away), and not couples that may be a part of a group; our assumption is that dyadic interaction and group interaction should first be studied separately. We take the distance when our bot first spots the avatars in range, otherwise our bot might affect the interaction, and it would no longer constitute a dyadic interaction.

Also, unlike Yee et al. [18], we do not rely on the chat texts to judge whether the avatars are talking, because users could be communicating via a private instant-messaging channel, hidden from our bot.

4.2 Results

Our bot ran a few days and nights and collected 205 samples of pairs of avatars. Based on the criteria mentioned above we consider 49 of these couples to be interacting. An analysis of the distances reveals an interesting pattern (Figure 1): the distance between two avatars seems to have two distinct peaks: one around 1-1.5 meters and the other around 4.5 meters. Figure 3 shows a couple of avatars in these two distances. As a comparison, if we look at the distances of the 156 pairs of avatars assumed not to be interacting, the distances seem to be uniformly distributed (see Figure 2).

The precise interpretation of these results is not clear. Possibly, avatars first communicate when they are close enough to draw attention to each other (around 4-5 meters in SL), and if they feel close they move to a more personal distance (1-1.5 meters in SL). In any case, the important result is that the proximity among interacting avatars does not seem arbitrary. The null hypothesis that the samples are taken from a Gamma distribution is not rejected (Chi-squared = 36.9 on 48 degrees of freedom). In order to evaluate the fit we treat the distance (y) as the response variable with a Gamma distribution in a generalized linear model (GLIM), and then fit a constant as the independent variable. In such a case, if the fit was not good the deviance (which is the chi-squared value) would be high (at least 2 * the degrees of freedom); in this case the fit is good. This fit suggests that the distances are drawn from a random population with mean 2.85 and standard deviation 2.2 (although manual inspection suggests it may be a mix of two Gamma distributions). This implies that the distances follow the pattern of distribution of points randomly in space according to a Poisson process, i.e., there is some regularity in the spatial social behavior of avatars in SL, not unlike what we would expect in a real life social setting, such as a cocktail party.




Fig. 1. The distribution of the number of pairs of interacting avatars by distance, in half meter units.

We have tested whether there is a difference in proxemity related to (virtual) gender; Yee et al. [18] found such differences among male-male, female-female, and mixed dyads. As SL does not provide the gender of an avatar (neither virtual nor real), this information was not available to our bot. We deduced the gender of 44 out of the 49 pairs based on the avatar’s first name⁷; in our case there was no significant effect of gender on distance. As Yee et al. [18] mention, the evidence for such a gender effect in real life is mixed.

5 Spatial Response: An Automated Experiment

In real life, when two people are engaged in conversation, they typically respond to each other’s “body language”. One of the common effects is posture shift; when one person changes posture, the other would typically respond by changing their posture as well. If both are in rapport they would often mimic each other’s

⁷ In the future we recommend recording the display on the screen, for such post-hoc analysis.




Fig. 2. The distribution of the number of pairs of avatars who are not interacting with each other, by distance, in half meter units.

posture, but even if not there would almost always be a posture shift within a few seconds.

As mentioned earlier, posture is meaningless in SL, and we cannot test whether there is a posture-shift effect. Instead, our hypothesis is that if our bot would move towards the user it is interacting with, then the other user would also move in response. We do not predict what type of movement will take place; we merely predict there will be a movement response.

5.1 Method

In order to test this hypothesis we used our bot to carry out an automated experiment. We have programmed it to look for spatially isolated avatars and approach them. Based on the previous section, our instructions to our bot were as follows: stand 4 meters and face the other avatar. Then greet them (e.g., say “hello *name*, how are things?”), wait 2 seconds, and move towards them to a personal distance, of 1.2 meters. During the 10 seconds following this approach the bot observes whether the other avatar moves, and sends us the results.



(a)



(b)

Fig. 3. (a) Screenshots of two avatars interacting: (a) social distance — approximately 5 meters, and (b) personal space — approximately 1 meter.

It is possible, of course, to carry out such experiments with an experimenter controlling the avatar. One advantage of an automated experiment is accuracy and less experimental noise. Another advantage is that the researcher can spend time on the beach while the bot performs the experiment... It is still recommended to watch the screen (or record it and watch afterwards). For example, in one case our bot approached an avatar that was sitting inside a vehicle; naturally there was no spatial reaction from that avatar.

First, in order to establish a baseline, We have analyzed the data obtained from 49 interacting pairs of avatars (in Section 4) to find out how often avatars change position during interaction. We found out that the rate of movement is very slow, and many couples did not move at all. On average, we have observed 65 position changes in 172 minutes of interaction logged, or an average of one motion every 2 minutes and 39 seconds.

Next, we sent our bot to carry out the experiment, in 10 different locations selected arbitrarily. Our bot approached 28 avatars, and sampled for avatar motion in the 10 seconds after approaching them.

5.2 Results

The results indicated that 12 out of the 28 opponent avatars moved within the 10 seconds duration following the approach. On the one hand, this is not a very reliable response; it means that only 42% of the users responded as we expected them to⁸. On the other hand, this is significantly more than position shifts in typical SL interactions: based on the baseline, we could expect 1 or 2 motions (280 seconds divided by 159 = 1.76) overall; we found the response to be 6 times larger than the baseline, so we regard this as an indication that many of the users did respond to our bot's approach.

As a side note, our bot recorded the contents of the chat channel after greeting the avatars, and sent us the data. 20 out of the 28 users responded to our bot's greetings ("Hello *name*, how are you doing?"). These were always short casual responses, mostly "hi", or slightly more sophisticated response such as "hmm... doing what?". We could see this as an indication that the majority of people respond to the verbal (chat) channel, more so than to the non-verbal, spatial channel.

6 Discussion and Future Work

We have described how our SL bots can systematically collect data from SL, and even carry out automated experiments. Ideally, we would like our SL bots to be able to socialize successfully in SL. Currently they are able to approach single avatars and, with a high rate of success, illicit some response (either verbal or position change). Such studies are, in general, useful for designing avatars and autonomous virtual agents with intelligent spatial behavior.

⁸ We did not find equivalent numeric data regarding the frequency of posture shifts in real life, but we assume the response is close to a 100%.

We have used scripting language (LSL) for our study, and it provided for most of our needs. The SL client have been recently released as open source, and there is an intention on behalf of Linden Labs to release the server source code as well. Note that this would allow more flexibility in implementing the agents, but would not address our main needs. First, we would still not have access to statistics available on Linden servers (with its very large number of users). Second, it is not clear if the information we need is explicitly available in the SL system. For example, we would like to know what avatars look like (e.g., are they male, female, or non-human). Using the client source one can extract scene information, including both the 3D model or the 2D rendered view. However, we would still require significant (virtual) machine vision effort to be able to determine what the avatar looks like, e.g., whether the avatar is attractive or not.

Furthermore, if we want our bots to be able to engage in more meaningful interactions, we hit the natural-language dialogue barrier. This is notoriously difficult, but we suggest SL may be an interesting domain for studying natural-language interaction. Unlike online chat environment, the conversation is often situated; for example, a user may tell another user: “do you see that green sign above our heads? click it to participate in the lottery.” To further study non-verbal mediated communication we need to stay in the realm of highly-immersive VR settings, including at least some information on postures and gestures, extracted from trackers.

Bots collecting information, and even carrying out experiments, raise methodological issues. For example, we need to find ways to ensure that we are sampling SL properly. We need to find ways to know who are the users behind the avatars we are interacting with: are they currently socializing, or are they working? in the future we might need to make sure they are not bots themselves...

Such research may also raise ethical, and even legal considerations. Carrying experiments with human subjects generally requires following ethical guidelines and experiments require formal approval from an official committee. For example, informed consent is always required from a subject before participating in an experiment. As long as the study is purely observational there is no problem. In our case the manipulation of the subjects (in Section 5) is clearly insignificant. If there were direct information being asked about the subjects, then it may be necessary to get institutional ethics approval, and also to make it clear to the subjects that this is a study, and they have the right to refuse, as in any real-life experiment. There is still a a difference from real-life experiments, since everyone is anonymous. However, some people put enormous resources into maintaining online personas, and would rightly object to using that persona’s name in records. Another concern will need to be addressed if the bots would be able to carry out more meaningful social interactions, since this would involve deception. In some virtual environments users are happy to accommodate bots in the condition that they declare themselves to be bots.

We have found some evidence of spatial social behavior in SL. Although these are early results, we see that: i) users tend to keep their avatars in non-arbitrary

proximity from the other avatars they are interacting with, and ii) almost half of the users responded in to an approach by our bot into their personal space by changing their position.

We agree with Yee et al. [18] that even in such a non-immersive experience as SL (small avatars observed on a computer screen, controlled by mouse and keyboard), users assign some social significance to the spatial context of their avatars. However, we do not agree with their conclusion that spatial social behavior in SL is identical to real-world behavior. Rather, we suspect that such virtual proxemics is a result of a transformation of real life proxemics. For example, while psychological literature bounds interaction distance by 12 feet, or 3.66 meters [6, 7], we found a significant number of avatars interacting with distances of up to 5 (virtual meters); such differences would not be surprising, since SL lacks the richness of the NVC channel, and relies of text chat rather than speech. Also, these differences could be due to a small field of view in SL, as compared with real life.

As another example, SL avatars rarely move while interacting, while we expect people in real life to change positions much more frequently within an interaction. These pieces of evidence testify to the possibility that SL users do not have a high sense of presence.

Nevertheless, we still see it important to study social behavior in non-immersive environments such as SL. People are now spending a lot of time in such environments (not only in SL but more so in multi-user video games), and we suggest there is both theoretical and practical reasons for comparing people's behavior in such environments with their behavior in real life, and also with their behavior in highly-immersive VR. Practically, such research will inform us how to construct improved interfaces, which preserve more of the social nuances of real-life interactions. Theoretically, such research may provide us with new insights about the human brain and psychology, from the way they adapt to new types of experiences.

ACKNOWLEDGEMENTS

This work has been supported by the European Union FET project PRESENCIA, IST-2006-27731.

References

1. J. N. Bailenson, J. Blascovich, A. C. Beall, and J. M. Loomis. Interpersonal distance in immersive virtual environments. *Personality and Social Psychology Bulletin*, 29:1–15, 2003.
2. S. Benford, J. Bowers, L. E. Fahlen, C. Greenhalgh, and D. Snowdon. User embodiment in collaborative virtual environments. In *Proc. of CHI'95*, pages 242–249, 1995.
3. J. Bowers, J. Pycock, and J. O'Brian. Talk and embodiment in collaborative virtual environments. In *Proc. of CHI'96*, pages 58–65, 1996.

4. W. Gibson. *Neuromancer*. Voyager, 1984.
5. E. T. Hall. *The Silent Language*. New York: Doubleday, 1959.
6. E. T. Hall. *The Hidden Dimension*. New York: Doubleday, 1966.
7. L. Hayduk. Personal space: Where we now stand. *Psychological Bulletin*, 94:293–335, 1983.
8. I. Heldal, R. Schroeder, A. Steed, A. S. Axelsson, M. Spante, and J. J. Widestr?m. Immersiveness and symmetry in copresent scenarios. In *IEEE Virtual Reality 2005*, pages 171–178, 2005.
9. G. A. Kaminka, M. M. Veloso, S. Schaffer, C. Sollitto, R. Adobbati, A. N. Marshall, A. Scholer, and S. Tejada. Gamebots: A flexible test bed for multiagent team research. *Commun. ACM*, 45(1):43–45, 2002.
10. J. E. Laird. Research in human-level AI using computer games. *Commun. ACM*, 45(1):32–35, 2002.
11. J. M. Loomis, J. J. Blascovich, and A. C. Beall. Immersive virtual environments as a basic research tool in psychology. *Behavior Research Methods, Instruments, and Computers*, 31(4):557–564, 1999.
12. N. Stephenson. *Snowcrash*. ROC Publishing, 1991.
13. J. G. Tromp, A. Steed, and J. R. Wilson. Systematic usability evaluation and design issues for collaborative virtual environments. *Presence: Teleoper. Virtual Environ.*, 12(3):241–267, 2003.
14. H.H. Vilhjlmsson and J. Cassell. Bodychat: Autonomous communicative behaviors in avatars. In *Int'l Conf. Autonomous Agents*, pages 269–276, Minneapolis, Minnesota, 1998.
15. V. Vinayamoorthy, A. Brogni, A. Steed, and M. Slater. The role of posture in the communication of affect in immersive virtual environments. In *The 2nd ACM SIGGRAPH International Conference on Virtual Reality Continuum and Its Applications*, 2006.
16. V. Vinayamoorthy, A. Steed, and M. Slater. The impact of a character posture model on the communication of affect in an immersive virtual environment. 2006. submitted.
17. M.V. Sanchez Vives and M. Slater. From presence to consciousness through virtual reality. *Nature Reviews Neuroscience*, 6(4):332–339, 2005.
18. N. Yee, J. N. Bailenson, M. Urbanek, F. Chang, and D. Merget. The unbearable likeness of being digital; the persistence of nonverbal social norms in online virtual environments. *Cyberpsychology and Behavior*, 10:115–121, 2007.