The Effect of Avatar Connectedness on Task Performance

John M. Linebarger and G. Drew Kessler

Technical Report LU-CSE-02-007
Lehigh University, Department of Computer Science and Engineering
{jmlg,gdk2}@lehigh.edu

Abstract

Previous research has examined the relationship between avatar representation and the sense of presence and co-presence in a shared virtual environment. A positive correlation has been found between the realism of the representation and both categories of presence. However, to our knowledge no research has focused explicitly on the effect of avatar representation in general and connectedness or embodiment in particular on task performance in a virtual environment. Three experiments are described which directly address this issue. Two of the experiments involve the representation of the self avatar in an immersive virtual environment in performing repetitive, semi-precise and non-repetitive, precise tasks. The third addresses the representation of avatars of other users, in both immersive and desktop virtual environment interfaces, in performing direct and avatar-mediated object-focused tasks. Four avatar representations were tested, each of which was either "connected" (i.e., embodied) or not, and "correlated" (i.e., color coded) or not. No significant difference in task completion time was observed between comparable self avatar representations in either task category, or between the representations of avatars of other users in direct object-focused tasks. For avatar-mediated object-focused tasks the representation was significant, with correlation having a much greater impact than connection on task completion times. Thus simpler, less computationally expensive avatar representations are quite adequate for task performance in a virtual environment for certain kinds of tasks.

1. Introduction

Several recent studies have investigated the relationship between the realism of the avatar representation and the sense of presence and co-presence in the virtual environment; for example, see [1] and [2]. In general, presence and co-presence have been shown to be positively correlated with the realism of the avatar representation. And "realistic" has generally implied "embodied," or at least "connected" or "articulated."

But how important are presence and co-presence when the purpose of the shared virtual environment is to perform a task? We are investigating the support of group collaboration modes in a shared virtual environment for product design and wish to ensure that the avatar representation chosen is adequate for the shared task. In preparation for the investigation we are developing a framework for distributed virtual environments called Simple Shared Virtual Environment (SSVE), as well as several applications that use the framework. One of them, a Chinese Checkers game intended to train virtual teams in the use of shared virtual environments, tracks only the positions of the head and one hand and utilizes an extremely simple avatar representation—a head mounted display (HMD) icon and a stylus. Any confusion about which stylus belongs to which HMD icon is resolved by attaching a color-coded marble to the end of the stylus to indicate the color of the player. This simple expedient has proved quite successful in practice. So might disconnected (non-embodied) avatar representations be adequate for certain kinds of tasks?

2. Motivation and Previous Work

We were motivated by both pragmatic and theoretical concerns. The immersive virtual environment resources at the Graphics, Interaction, and Virtual Environments (GIVE) lab [3] at Lehigh University are relatively modest, consisting of only two HMDs and two four-port tracking devices. Sophisticated avatar representation schemes, such as those in the seventeen-sensor version developed at MIRALab and Computer Graphics Lab, EPFL [4], are thus out of the question. Even if only three or four sensors were involved and heuristics were used to calculate the remaining joint positions, we wondered whether the computational and view occlusion costs would be worth the performance benefits for certain tasks. On a more conceptual level, we were influenced by the work of Hindmarsh *et al.* on object-focused interaction [5] but questioned the necessity of the choice of an embodied avatar representation. In particular, we thought that a category of distinction for object-focused tasks should be introduced—*mediation*—and that avatar-mediated object-focused tasks should be distinguished from direct object-focused tasks. In avatar-mediated object focused tasks, the focus has to pass through the avatar to identify the object of interest; in direct object-focused tasks the object of interest can be located directly. (Table 1 provides examples of the two categories.) Since collaborative product design is indeed an object-focused task, the research question became, "Might a disconnected avatar representation be appropriate for some kinds of object-focused tasks?" Our operating hypothesis was that the avatar representation chosen would *not* matter for direct object-focused tasks, but *would* matter for avatar-mediated object-focused tasks.

Direct	Putting one box on top of another Mirroring a marble movement Determining a legal move in a board game
Avatar-mediated	Deictic references (pointing) Following a gaze Proximity determination (*e.g.*, the closest object to an avatar)

Table 1: Examples of object-focused tasks by category

For completeness, we addressed the self avatar representation as well. Our operating hypothesis here was that a connected representation would lead to better task performance because of the power of proprioception; see [6] for a classic discussion of this issue. We also were interested in comparing task performance in immersive and desktop environments. This was driven by practical concerns—given our laboratory configuration, a six-person collaborative product design team would consist of two immersed and four desktop participants, and we wished to quantify the performance implications of each interface along with the avatar representations.

To the best of our knowledge, no previous research has directly addressed the issue of the relationship between avatar representation, particularly avatar connectedness or embodiment, and task performance. As stated above, much of the previous research has investigated the relationship between avatar realism and the sense of presence and co-presence. A small amount of research has focused on the relationship between presence and task performance; however, that the relationship was positive has largely been assumed rather than demonstrated (see [7]). Indeed, some studies find that presence can have a negative effect on performance; see [8] and [9] for recent examples. Still, any assumption that avatar realism positively affects task performance seems to flow through presence; we call this the "task performance transitivity assumption." Our goal in this series of experiments was to investigate the relationship between avatar connectedness and task performance directly.

3. Experimental Methodology and Procedure

Two experiments were originally conceived, one dealing with self avatar representations and one with avatar representations of other users. A third, self avatar, experiment was conducted to clarify results of the first self avatar experiment. A within-subjects design (also known as a randomized complete block design) was chosen in order to increase the power of the results. Time in seconds and number of errors were the primary response variables for each of the experiments. Avatar connection and correlation were the primary factors; each had two factor levels. In the "other avatars" experiment, the interface (immersive or desktop) was also a factor, and the three factors were completely crossed. A series of randomized training runs was conducted before each experiment.

The subjects volunteered for the experiment primarily in response to broadcast email solicitations, but also in response to friends. There were 20 participants for each experiment for a total of 60 subjects. 46 of the subjects were male, and 14 were female. Many were graduate students in their twenties, but the ages ranged from 18 to 60. None identified themselves as being color-blind; only two were left-handed. Nine subjects participated in more than one experiment.

Each subject was asked to sign a consent form and was given both a verbal and a written description of the experiment. To guard against the possibility of simulator sickness in the immersive environment, the subject was asked to complete a slightly modified version of the Simulator Sickness Questionnaire (SSQ) [10], both before and after the experiment, as a screening and evaluation instrument, respectively. If any symptom was rated as moderate or above, the person was asked either to return later (pre-experiment) or to remain in the lab until the symptom subsided (post-experiment). Only three such symptoms were observed, two pre-experiment and one post-experiment. At the conclusion of the experiment the subject was asked for their qualitative assessment of the avatar representations. The investigator's questions generally took the form of, "Which avatar representation did you prefer for the task, if any, and why? And which interface did you prefer, if any, and why?"

The software for the experiment was written using the SSVE distributed virtual environment framework, which is built on top of the Simple Virtual Environment (SVE) library [11]. The immersive versions of the experiment were rendered in frame-interlaced stereo. A stripped-down version of the Chinese Checkers application mentioned above was used as a base for the considerable custom C++ coding required for each experiment. STATGRAPHICS Plus version 5.0 Enterprise Edition was used to calculate the analysis of variance (ANOVA) at a 95% confidence level. The hardware configuration consisted of a four-port Polhemus 3SPACE FASTRAK electromagnetic tracking device with a one-button stylus on the first port; a Virtual Research Systems V8 head mounted display with 640x480 resolution and a 60 degree field of view; a Dell Precision 330 workstation with a 1.5G Intel Pentium 4 processor, 512MB of RAM, a two-button mouse with a scroll wheel, and two 3Dlabs Wildcat II 5110 graphics cards, running Windows 2000 with Service Pack 3; and two 21-inch Dell P1100 UltraScan monitors with Trinitron tubes. To reduce the number of experimental factors, the same investigator and hardware platform were used for all trials.

4. Avatar Representations

As Table 2 indicates, the two crossed avatar factors, connection and correlation (implemented here as embodiment and color-coding), yielded four different avatar representations. The unconnected representations simply reflected the position and orientation information from the trackers using an appropriate icon. To generate the connected representations, a "widget" was developed using the SVE 3D interactor framework [12] that creates embodied humanoid avatars (from the waist up) using only two or three tracking sensors. If two sensors were used, the head, torso, and right arm were rendered; if three sensors were used the left arm was depicted as well.

A brief description of the connected avatar generation algorithm is in order. The avatar is constructed from a set of rigid body sections positioned on the basis of the tracker data, a set of measurements for the user (with reasonably general defaults), and a set of heuristics. The positions tracked are assumed to be the top of the user's head and a location in the center of a closed hand. The objects rendered are a head (with neck), a torso, an upper arm, a lower arm, and a hand (with wrist). The head object position and orientation are based on the head tracker (the top of the head). The torso begins a given distance "below" the tracked head in the same coordinate system. The height of the base of the neck is compared to a given user height to determine the forward lean of the torso; it is assumed that the legs are not bent and are not leaning in any direction. The base of the wrist is determined by the tracked hand sensor. If that position is at or lower than the height calculated by subtracting the length of the arm from the base of the neck, the hand is assumed to be at the user's side and the torso is twisted so that its side is towards the hand. If the height of the wrist is at or above the height of the neck, it is assumed that the user is reaching forward, and the torso is twisted so that the hand is forward of it. The upper and lower arm segments are either straight (if the hand tracker is farther from the base of the neck than the given arm length) or bent such that the elbow is downward and to the right.

Due to experiment time constraints, the default user measurement settings were used instead of being customized for each user. However, on occasion the shoulder width setting of the self avatar representation was adjusted to prevent the shoulder from occluding the view.

	Correlated	**Uncorrelated**
Connected	Humanoid upper body using three-sensor heuristics Color-coded wrists	Humanoid upper body using three-sensor heuristics Body-colored wrists
Unconnected	HMD icon (head) Color-coded stylus (hand)	HMD icon (head) Bare stylus (hand)

Table 2: Avatar representation matrix

5. Self Avatar Experiment

The initial self avatar representation experiment consisted of two tests, given in random order. One test was a repetitive semi-precise task in which the subject raced a marble hole-by-hole from one end of the Chinese Checkers board to the other. At least sixteen discrete marble selections, movements, and releases were required to accomplish this task using legal Chinese Checkers moves. The stylus device was used to select the marble by placing the tip of the representation just inside the marble and pressing the button. A successful selection was indicated by an audible "whoosh" and a change in marble color to gray, thus substituting two senses for the loss of the sense of touch. The marble was then attached to the stylus. After moving the marble to the desired spot the button was released, and the marble would audibly and visibly click into place on the nearest hole on the board. The task was considered semi-precise because the placement did not need to be exact, just closest to the target. The goal was to run the race in the shortest time with the fewest errors. Both the training set and the trials consisted of four races, one for each avatar representation.

The second test was a non-repetitive precise task, which consisted of placing a marble in a box elevated above the board. The width of the box in each dimension was identical to the diameter of the marble. The walls of the box were hollow and the laws of gravity did not apply, so the marble would remain at the spot where it was released. Five replications with each avatar representation were performed, in random order, for a total of twenty trials. The goal was to place all of the marble in the box; though rare, this goal was occasionally realized. Time was not important in this test. Note that the correlated (color-coded) versions of the avatars were used only to provide a comparison with the next experiment; it was anticipated that the results of the correlated and uncorrelated versions of an avatar would be essentially identical. Figures 1 and 2 show screen images of two different avatar representations used in the marble-in-box task.

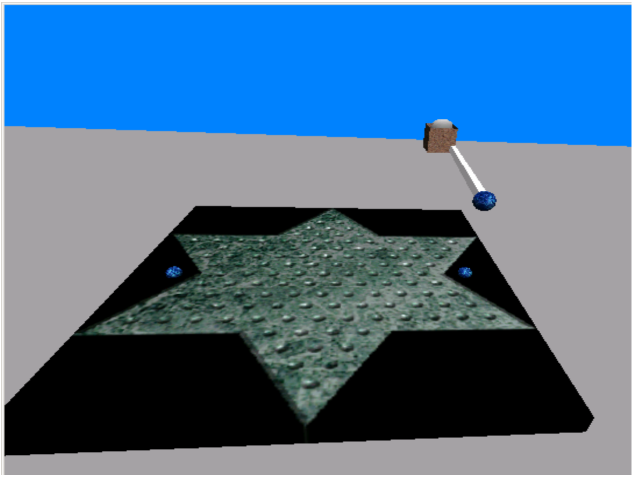

Figure 1: Unconnected, correlated self avatar

The investigator manually kept track of the number of errors and timed the marble race using a stopwatch, while the virtual environment program calculated the percentage of the marble in the box and the ratio of successful selections to attempted selections (which occurred when the user pressed the button on the stylus). The orders in which the four avatar representations appeared in both the training runs and the real trials were determined randomly.

In the post-experiment interview, 65% of the subjects said they preferred the hand representation for the race task and 70% preferred the stylus for the marble-in-box task. For each task, 10% expressed no preference. As shown in Table 3, the results of the race task reflect the preference—not only was avatar connectedness a significant factor in the race times recorded ($F_{1,19} = 14.65$), the two connected representations had a 12% lower average time than the two unconnected representations. Connectedness was a slightly more significant factor ($F_{1,19} = 15.13$) in the selection success ratio, providing a 7% gain. The avatar representation did *not* prove to be a significant factor for the number of errors.

Note that all table results are individually rounded to two or three decimal places. Significantly different pairs of means are marked with asterisks next to the second mean of the pair—a single asterisk for a p-value less than 0.005, a double asterisk for a p-value less than 0.0001. All unmentioned p-values, including interactions, were > 0.05.

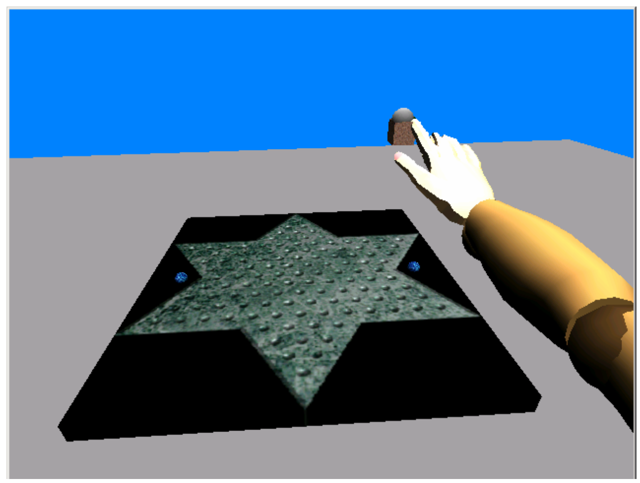

Figure 2: Connected, uncorrelated self avatar

Task	Response Variable	Avatar	Correlated	Uncorrelated	Mean
Race	Time (in seconds)	**Connected**	22.15	22.6	22.38
		Unconnected	25.2	25.4	25.3*
		Mean	23.68	24	
	Selection Success Ratio	**Connected**	0.831	0.82	0.826
		Unconnected	0.746	0.739	0.743*
		Mean	0.789	0.78	
	Number of Errors	**Connected**	0.6	1.05	0.825
		Unconnected	0.85	0.8	0.825
		Mean	0.725	0.925	
Box	Percent in Box	**Connected**	88.71	86.42	87.56
		Unconnected	86.89	87.16	87.02
		Mean	87.8	86.79	

Table 3: Table of means for the self avatar experiment (*p < 0.005)

However, for the marble-in-box task, the expressed preference for the stylus was not borne out by the results—the avatar representation did not significantly affect the percentage of the marble placed in the box. A "recoil" effect was frequently observed in this task, in which the marble would suddenly shift position when deselected. Post-experiment interviews identified the lack of an opposing force when pressing and releasing the stylus button as the possible cause. As a result, we would encourage device manufacturers to utilize opposing forces for press and release to ensure placement accuracy for precise tasks.

6. Avatars of Other Users Experiment

The "other avatar" representation experiment also consisted of two tasks. However, each task was conducted twice, once in an immersive environment and once in a desktop environment. The first task was designed as a direct object-focused task. The goal was to mirror the marble movement of an avatar on the other side of a Chinese Checkers board. The marble movement was randomly chosen from a set of twelve possible moves from the starting position, of which half were one-hole moves and half were two-hole moves. Eight training runs were followed by twenty-four real trials, which meant that six replications of each avatar representation were presented in random order. Three of the six replications were for one-hole moves, the other three for two-hole moves. The immersive version of the task used the stylus device as before, with the connected self avatar representation that was preferred in the first experiment. The desktop version used the mouse as the marble selection and movement mechanism. To select a marble, the subject placed the mouse cursor over the marble on the screen and pressed the left mouse button. A successful selection was indicated by an audible "whoosh" and a change in marble color to gray, and the marble was raised up to lie on a plane a short distance above the board. While the left mouse button was held down, mouse motion would propel the marble across the plane while keeping it directly underneath the cursor on the screen. A shadow was placed on the board under the marble to aid in depth perception. When the marble was over the desired hole, the subject released the left mouse button and the marble clicked into place. Figure 3 depicts this task with a connected, correlated other avatar.

The second task was designed as an avatar-mediated object-focused task. Five avatars were arrayed around a Chinese Checkers board. Each of them selected the marble closest to them and moved it one hole away on the board. The goal of the task was to determine if any of the avatars had touched a marble of the wrong color. The information the subject could use to answer the question was always present both on top and in front of each avatar's head. For correlated representations, it was present at each avatar's hand as well. This task also consisted of eight training runs and twenty-four real trials and thus of six replications per avatar representation. For three of the replications a randomly chosen avatar selected a marble of the wrong color; for the other three, each avatar picked the correct marble. The subject was told the set of possible colors in advance—red, green, cyan, purple, white, and yellow. Figure 4 portrays this task with an unconnected, correlated avatar representation.

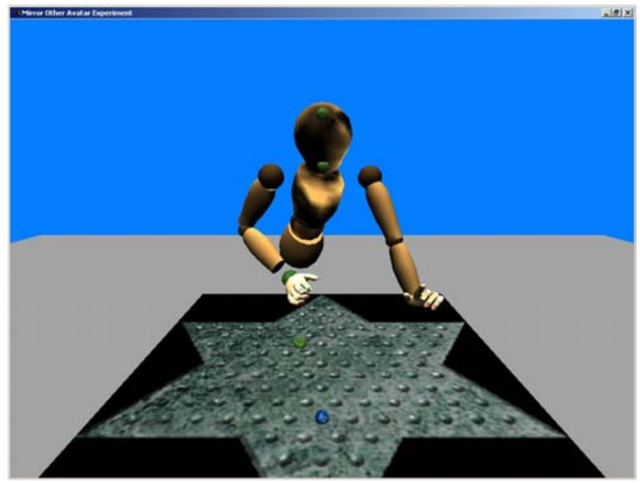

Figure 3: Mirror marble movement of another avatar task

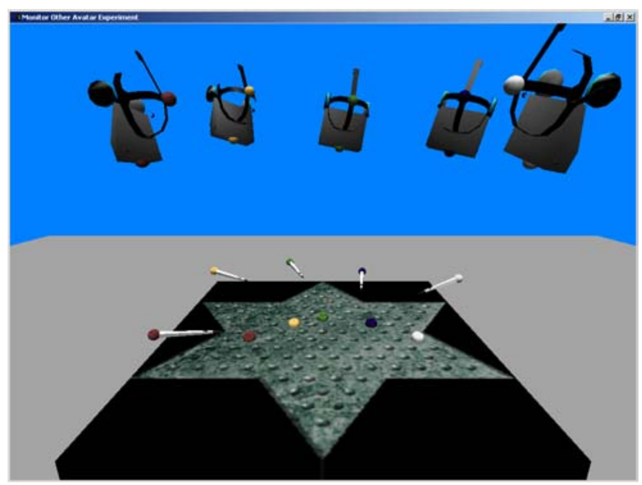

Figure 4: Monitor marble selections of other avatars task

Properly randomizing this experiment proved so challenging that a custom program was written which determined the order of the four tasks, the training runs for each task, and the actual trials. The printout from that program guided the investigator in the operation of the experiment. In addition, the virtual environment program itself made the timing measurements. For the mirror marble movement task, the clock started automatically when the other avatar selected the marble and stopped with a key press when the subject indicated that the marble movement was successfully mirrored. For the avatar marble selection monitoring task, the clock started when the scripts began playback and ended with a key press as soon as the subject verbalized a conclusion. To ensure that the avatar movements were realistic and repeatable, a custom SVE 3D interactor widget was created that recorded object position changes over time to a file and played them back when attached to objects in the scene. Up to four files were needed for each avatar, two for the hands, one for the head, and one for the marble. The position scripts were recorded separately in advance. Each was approximately twenty seconds long, and the marble was selected about three seconds into the script.

In the post-experiment interview, 50% of the subjects expressed a preference for a connected representation in the marble mirroring task, primarily because they felt it better communicated the direction and intention of the marble movement. However, 35% had no preference, generally because their focus was on the marble and not on the avatar, and no preference was reflected in the results (shown in Table 4) for either time or error.

Response Variable	Interface	Avatar	Correlated	Uncorrelated	Mean
Time (seconds)	Immersive	**Connected**	5.237	5.56	5.399
		Unconnected	5.169	5.432	5.301
		Mean	5.203	5.496	5.35
	Desktop	**Connected**	4.422	4.385	4.403
		Unconnected	4.355	4.334	4.34
		Mean	4.388	4.36	4.374**
Num. Errors	Immersive	**Connected**	0.008	0.05	0.029
		Unconnected	0.057	0.042	0.049
		Mean	0.033	0.046	0.039
	Desktop	**Connected**	0.062	0.042	0.052
		Unconnected	0.058	0.017	0.038
		Mean	0.06	0.029	0.045

Table 4: Table of means for the task of mirroring the marble movement of another avatar (**p < 0.0001)

Other than the particular subject, what *did* significantly affect the marble mirroring task? The interface used was not significant with respect to error, but was highly significant with respect to time ($F_{1,803}$ = 104.95). The desktop interface was the clear winner overall, with a mean time almost a second (18%) less than that of the immersive interface. Interactions were also examined, with only the subject-interface interaction showing any significance ($F_{19,803}$ = 7.24, p-value < 0.0001). This indicates that subjects had an affinity for a particular interface, which was somewhat surprising considering that only half the subjects expressed an interface preference, which was evenly split.

For the task which monitored the marble selections of avatars of other users, the post-experiment interview indicated an overwhelming preference for the color-coded stylus representation (80%), and 35% of the subjects rated the non-color-coded stylus the worst. Several subjects stated that they found the connected avatars "distracting" or "too much to look at," and that they often got in the way of the view. Only a slight preference for the desktop interface was expressed (20% to 10%), primarily because of its wider field of view. This time the subjective preferences were clearly reflected in the faster times that were recorded. Table 5 below summarizes the results of this task. Avatar correlation was highly significant for both correct and incorrect selections ($F_{1,320}$ = 255.08 and $F_{1,321}$ = 42.82, respectively). Avatar connection was significant only for the correct selection case ($F_{1,320}$ = 10.11). Significant (but to a lesser extent than all but connection) interactions were found between connection and correlation for both correct and incorrect selections ($F_{1,320}$ = 46.98 and $F_{1,321}$ = 16.16), and between interface and correlation for correct selections. The unconnected, correlated representation always had the lowest mean time, and in all cases but one the unconnected, uncorrelated representation had the highest. Note that the data for this task are broken out by correct and incorrect selections; if all of the avatars selected the correct marble the task essentially became an exhaustive search, which we expected to take longer on average than a search that could be truncated by an incorrect marble selection.

The results for the error response contained something of a surprise. Avatar correlation had a significant effect ($F_{1,321}$ = 8.51, p-value = 0.0038) but only for the incorrect selection case. Further investigation shows that the highest error rate was produced by the unconnected, correlated representation—which is also the representation associated with the lowest time scores—especially on the desktop. Why might that have happened? One guess is the "snap judgment" phenomenon; the average response time for that avatar representation is over a half second (14%) less than the average of the mean response times of the other representations for incorrect selections. Perhaps the uncluttered representation, the wider field of view of the desktop interface, and the shortcut allowed by the color correlation fostered hasty conclusions. Anecdotal evidence points to this possibility as well—several times the investigator heard someone come to a quick conclusion and then say "No, wait!" However, another possibility is that using color as the correlation mechanism could lead to more false negatives than other possible correlation mechanisms, such as shape correspondence or textual labeling.

Response Variable	Correct Selection	Interface	Avatar	Correlated	Uncorrelated	Mean
Time (seconds)	Yes	Immersive	**Connected**	4.314	5.037	4.675
			Unconnected	3.466	5.626	4.546*
			Mean	3.89	5.331**	4.611
		Desktop	**Connected**	3.612	4.161	3.887
			Unconnected	3.06	4.084	3.572*
			Mean	3.336	4.123**	3.729**
	No	Immersive	**Connected**	4.059	4.28	4.17
			Unconnected	3.383	4.964	4.174
			Mean	3.721	4.622**	4.172
		Desktop	**Connected**	3.435	3.796	3.615
			Unconnected	3.24	4.094	3.667
			Mean	3.338	3.945**	3.641**
Num. Errors	Yes	Immersive	**Connected**	0.017	0.017	0.017
			Unconnected	0.0	0.017	0.008
			Mean	0.008	0.017	0.013
		Desktop	**Connected**	0.0	0.0	0.0
			Unconnected	0.0	0.0	0.0
			Mean	0.0	0.0	0.0
	No	Immersive	**Connected**	0.117	0.05	0.08
			Unconnected	0.179	0.092	0.135
			Mean	0.148	0.071*	0.109
		Desktop	**Connected**	0.117	0.113	0.115
			Unconnected	0.275	0.083	0.179
			Mean	0.196	0.098*	0.147

Table 5: Table of means for the task of monitoring the marble selections of other avatars (*p < 0.005, **p < 0.0001)

But which of the two factors, avatar correlation or interface, had the greater effect? The F-ratios tell the tale; the ratio for the interface for both correct and incorrect selections ($F_{1,320}$ = 159.55 and $F_{1,321}$ = 21.22, respectively) was always lower than the corresponding ratio for avatar correlation. Again, it was the desktop interface that yielded the lower time scores, by an average of 16%.

7. Revised Self Avatar Race Experiment

After analyzing the results of the second experiment, we were tempted to conclude that an unconnected but correlated avatar representation should be adopted for the collaborative product design application. But the results of the first experiment favored the connected representation. Should two different types of avatar representations be chosen, one for the self avatar and one for avatars of other users? Reviewing the feedback from the first experiment, we found two possible explanations. The first was the phenomenon we expected, that the sense of the virtual arm being connected to the subject's real body proved helpful in performing the task of repetitive selection and movement. However, several subjects mentioned that the slender black tip of the stylus icon made it difficult to select the marble, especially when the stylus was some distance from the eyepoint, because perspective projection of the tip made the hot spot for selection difficult to see. Was it proprioception or resolution problems that led to better performance with the connected self avatar?

To disentangle these potentially confounded factors, we conducted a short final experiment, which was a modified version of the marble race task in the first experiment. The stylus icon in the unconnected representation was enlarged to be the same size as the index finger of the connected representation. In addition, the sharp black stylus tip was removed, the front of the stylus was tapered to resemble the shape of the index finger, and the tip was given the same color as the nail of the index finger. These were the only two representations tested; the color-coding was dispensed with. Data collection problems prevented us from reliably capturing the selection success ratio. All other details were identical to the first experiment.

The outcome of the post-experiment interview was quite similar to the first experiment, with 65% preferring the connected representation and 10% expressing no preference. However, this time the preference was not borne out in the data (shown in Table 6), either for the race time or the number of errors—the avatar representation was not a significant contributor to either of them.

Avatar	Time	Errors
Connected	22.16	0.875
Unconnected	23.05	0.7
Mean	22.6	0.788

Table 6: Table of means for the revised self avatar race

8. Conclusions and Future Work

Several conclusions can be drawn. The representation of comparable self avatars was not shown to have a significant effect on task performance, at least for certain kinds of tasks. The representation of the avatars of other users also had no significant effect on task performance for the direct object-focused task. But for the avatar-mediated object-focused task the avatar representation had a significant effect, although the correlation of the representation had a much larger impact than connection. For direct object-focused tasks it was the interface that had the greatest effect on task performance; for avatar-mediated object-focused tasks it was correlation. The unconnected, correlated representation had the best time scores, though its error rates were higher; the interface with the faster times was the desktop.

Future work includes the exploration of correlation mechanisms other than color, as well as the utility of a "lightweight" connected avatar consisting of a line drawn from the top of an HMD icon to the top of a stylus icon. In addition, based on the observation that for direct object-focused tasks the avatar representation could consist merely of a color-coded stylus, the value of representing the head at all could be questioned. Our hypothesis is that both the head and the hand communicate *attention*. A split in attention communicates *intention*, in which the head is assumed to be the leading indicator and the hand the lagging indicator. A future experiment is planned to test this hypothesis.

References

[1] Juan S. Casanueva and Edwin H. Blake. "The Effects of Avatars on Co-presence in a Collaborative Virtual Environment." Technical Report CS01-02-00, Department of Computer Science, University of Cape Town, South Africa, 2001.

[2] M. Gerhard, D.J. Moore, and D.J. Hobbs. "Continuous Presence in Collaborative Virtual Environments: Towards a Hybrid Avatar-Agent Model for User Representation." *Proceedings of the Third International Workshop on Intelligent Virtual Agents (IVA 2001),* 10-11 September 2001, Madrid, Spain: 137-155.

[3] http://www.cse.lehigh.edu/GIVE

[4] Tolga K. Çapin, Igor S. Pandzic, Nadia Magnenat-Thalmann, and Daniel Thalmann. *Avatars in Networked Virtual Environments*. Chichester, West Sussex, England: John Wiley & Sons, Ltd., 1999.

[5] Jon Hindmarsh, Mike Fraser, Christian Heath, Steve Benford, and Chris Greenhalgh. "Object-Focused Interaction in Collaborative Virtual Environments." *ACM Transactions on Computer-Human Interaction (TOCHI)* 7, No. 4 (December 2000): 477-509.

[6] Mel Slater and Martin Usoh. "Body Centered Interaction in Immersive Virtual Environments." In Magnenat Thalmann and Daniel Thallman, eds., *Artificial Life and Virtual Reality*, (Chichester, West Sussex, England: John Wiley and Sons Ltd., 1994): 125-147.

[7] Robert B. Welch. "How can we determine if the sense of presence affects task performance?" *Presence* 8, No. 5 (October 1999): 574-577.

[8] Grace M. Bochenek and James M. Ragusa. "Virtual Collaborative Design Environments: A Review, Issues, Some Research, and the Future." *Proceedings of the 2001 Portland International Conference on the Management of Engineering and Technology (PICMET '01)*, Vol. 2, 29 July-2 August 2001, Portland, Oregon, USA: 726-735.

[9] K. Mania and A. Chalmers. "A User-Centered Methodology for Investigating Presence and Task Performance." *Proceedings of the 3rd International Workshop on Presence (Presence 2000)*, CD-ROM, March 2000.

[10] Robert S. Kennedy, Norman E. Lane, Kevin S. Berbaum, and Michael G. Lilienthal. "Simulator Sickness Questionnaire: An Enhanced Method for Quantifying Simulator Sickness." *The International Journal of Aviation Psychology* 3, No. 3 (1993): 203-220.

[11] G. Drew Kessler, Doug A. Bowman, and Larry F. Hodges. "The Simple Virtual Environment Library: An Extensible Framework for Building VE Applications." *Presence* 9, No. 2 (April 2000): 187-208.

[12] G. Drew Kessler. "A Framework for Interactors in Immersive Virtual Environments." *Proceedings of 1999 IEEE Virtual Reality Conference (VR '99)*, 13-17 March 1999, Houston, TX, USA: 190-197.