

RehabJim: A Third Person Approach To Virtual Reality Biomechanical Rehabilitation

Abstract

RehabJim is a Unity3D application developed on behalf of researchers at a major Rehabilitation Institute in order to explore the opportunities that Virtual Reality may offer to biomechanical and neurological rehabilitation. In particular, our work focuses on arm actions performed while standing by patients post-stroke. We leverage a 3D immersive environment, augmented with Kinect interaction. Our approach follows a third person perspective of the patient's body, and employs two cartoon-style avatars — the patient and their therapist. The user is asked to reach to computer-generated virtual objects with her hands. Our evaluation with users and a domain expert shows that this type of environment leads to an engaging, enjoyable experience that can encourage patients to perform a wide variety of whole-body motions.

Keywords: rehabilitation avatars, virtual worlds, virtual reality, CAVE2

1 Introduction

Patients come to Centers for Physical Therapy and Rehabilitation from all over the globe, “often with catastrophic or unusual conditions and always with great hope for an improved life ahead” [1]. In the quest to assist these patients, Virtual Rehabilitation is becoming popular in clinical research thanks to the development and availability of new immersive technologies. For example, the rehabilitation of post-stroke patients frequently uses virtual reality, and sometimes employs a first-person Serious Games ap-

proach to make therapy more enjoyable for patients (see Section 2). In such approaches, the patient experiences the virtual environment in a first-person view.

Unlike the existing body of work, we present and evaluate an immersive virtual rehabilitation application which uses a third person-perspective: the patient controls with her body a 3D avatar. The application was commissioned by biomechanics researchers at a major center for rehabilitation research, and could be used to compare the learning-curve and time needed for rehabilitation across different therapies and environments. One of the main challenges in this work is the design of simple, non-distracting avatars and virtual world, which allows the patient to focus on performing correctly the rehabilitation exercises. An additional challenge comes from the tradeoffs typically associated with virtual reality, for example, loss of depth perception.

Our solution, RehabJim, is a third-person serious game for the rehabilitation of upper limbs in post-stroke patients, with particular emphasis on reaching movements. Our solution uses a state-of-the-art virtual reality environment, Unity 3D, and a Kinect controller (Fig. 1).

2 Related Work

Virtual Rehabilitation can be an effective way to provide rehabilitation to post-stroke patients. The literature reports on the strengths and opportunities represented by virtual rehabilitation [2]. One strength of virtual rehabilitation is that there is no difference in the perception of immersion in virtual reality environment between




Figure 1: External view of CAVE2 environment with a running instance of our application. The screens cover 320 degrees of visibility and a monitor at the entrance is used for managing the hardware resources.

healthy patients and post-stroke patients [3].

A possible way to approach virtual rehabilitation is through Serious Games. It has been demonstrated that this strategy is helpful to reduce the repetitiveness of exercises, which often implies a lack of motivation in patients [4]. An example of a serious games for rehabilitation is by Ma et al [5], in which patients perform several strengthening exercises for fingers, hands and arms using a high resolution HMD, Virtual Research 1280 stereo device. In the work of Petracca et al [6], a LEAP motion controller is used in combination with an adapted version of a web video game and the patient must keep her forearm in a fixed support. However, this type of setting is typically too constraining for the patient, while using a Kinect can give the patient freedom of movement.

Several attempts have been made using webcams. One example is the work of Hoermann et al [7], where a computerized mirror therapy was used: patients keep their hand(s) on a table in front of them. The hands are covered by a screen in which the mirrored unaffected hand is shown. Another example is the RGS work of Burke et al [8] where webcams, serious games and virtual rehabilitation are integrated. In this game, patients intercept spheres that are moving towards them. In this approach there is a clear advantage in that the therapy does not require a specialized environment. However, these games

can also be flat and not engaging. This might be a problem: a good virtual reality environment should be highly motivating [5], otherwise it is likely to have the same problems as traditional rehabilitation.

In the similar work of Mumford et al, patients need to move a physical object over an LCD horizontal panel [9]. This approach is not feasible for those patients who are unable to hold weights. Furthermore, our researcher collaborators suspect that moving objects might be distracting for patients [10] and thus could influence the patient performance.

The work closest related to ours is Turolla et al [11], where the users are also performing a reaching movement. Turolla et al also show the correct path on the screen, as we do. However, in their work the correct path is established by the therapist, while in our approach the path is computed automatically. In the work of Piron et al [12], which uses a form of tele-rehabilitation, the optimal path is also pre-recorded by the therapist. However, our collaborators argue that the presence of the therapist is an important aspect for rehabilitation and for this reason the therapist should also be present in the virtual world used for the rehabilitation.

A related technique often used in virtual rehabilitation is mirroring. For example, Patel et al [13] use virtual reality based mirrored feedback tasks. While we do not address mirroring in our framework, one of the available features is a distorted mode influenced by the mirroring technique.

3 Methods

We start with a technology and overall system overview, and then highlight our specific process and choices in the design of the avatar application. We then describe the main features for customizing the rehabilitation experience, and finally the user interface adopted for facilitating the interaction with the avatar and the system.

3.1 CAVE2 Environment

The CAVE2 [14] environment is a hybrid immersive reality system that combines the benefits of both scalable-resolution display walls

and virtual-reality systems. The technology allows information-rich analysis as well as virtual-reality simulation exploration. CAVE2 provides users with a 320-degree panoramic environment for displaying information at 37 Megapixels in 3D, approximately 24 feet in diameter and 8 feet tall and consisting of 72 near-seamless passive stereo LCD panels managed by a 36-node high-performance computer cluster. Thanks to the 10-camera optical tracking system, CAVE2 allows to know in real-time the position in space of a user who wears a particular pair of 3D glasses, and also the orientation of her head, guaranteeing a three-dimensional perspective that varies according to the movements of the user inside the space. At the same time, the system can track the 6DOF of multiple controllers, which can be used in order to interact with specific applications. CAVE2 additionally provides a 3D virtual sound system thanks to its multiple speakers.

In order to track the joints of the patients, we added a Microsoft Kinect v2 device to the CAVE2 environment. This device was placed at a height of two feet from the ground, with an inclination of 4 degrees. This arrangement allowed us to track the position over time of up to 29 human joints, and to calculate using matrix algebra their orientation in space. The maximum number of trackable users is currently 6, and the allowed user-distance from the device is approximately 12 feet. We also leveraged the Kinect v2 speech recognition feature in order to broaden the user interaction possibilities.

3.2 System Overview

Our application is developed in Unity 3D, and follows the constraints and the possibilities offered by CAVE2 environment. A third party wrapper allows us to define the position and orientation in space of each of the 64 screens, in order to render on each of them the correct portion of the scene. Each of the 32 computers then runs its own instance of the application and is controlled by a master node, which takes care of uploading content to the whole cluster and which makes synchronous requests to the other nodes. For example, generating an object in a random position would require the master node to compute the position and then transmit it to

all the nodes, otherwise each computer would produce a different result. Due to wire extension issues, the Kinect v2 device is connected to an external server which streams the position of the joints of the users to each node via a local network.

The application is typically initiated by a therapist from the master node of CAVE2, where simple data about the patient can be inserted for logging purposes. Then the patient can start her training session inside CAVE2 with or without the help of the therapist. Both patient and therapist are represented by an avatar in the virtual environment (Fig. 2) and their movements are tracked with the aid of the Kinect. Through the 3D user interface of the voice commands, each user can decide which training mode to start and which features to enable during the session. All data involving the training is then stored onto the master node inside a log file, which can be later used for post-training analysis.

3.3 Application Design

The application is meant to be engaging and enjoyable for patients, adopting the possibilities offered by virtual reality and leveraging the principles of gamification. Following several low-fidelity and high-fidelity prototypes and based on initial feedback from the rehabilitation researchers, we adopted a minimalistic environment characterized by few distinct natural features, a neutral background and an easily understandable color palette. Since an excessive number of elements would have provided a relevant source of distraction for some patients, we decided to keep the virtual scene as simple as possible.

The movements of the patient are tracked and reproduced by the virtual avatar moving in front of her. The domain experts considered photorealism a problem for the rehabilitation purpose — the small artifacts in the realistic avatars' appearance and behavior were distracting. We thus decided to adopt a stylized cartoon shape for our avatar, further maintaining coherency with the rest of the scene. The virtual objects that the patient has to reach are rendered as spheres, in order to have the distance from the user independent of the angle of approach.

A first issue regarding positioning virtual ob-



Figure 2: Part of the Tutorial Mode which presents the two type of avatars (therapist on the left, patient on the right).

jects involved excluding positions in which they could not be seen by the patient, considering that the avatar occupies a relevant portion of the screen. Since the avatar, having roughly human size, is virtually situated in front of the user and is seen from the back, objects need to appear even further than the avatar in order to be caught: for this reason, our custom seeding algorithm generates objects only in positions reachable in relation to the current position of the patient and never directly in front of her.

A second issue involves the perception of depth in the 3D virtual world: in virtual environments, sometimes, it is not easy to distinguish how far an object is and how much we should move in order to reach it. By experimenting with different alternatives, we decided to use the camera angle to place realistic shadows on the scene (Fof. 3 left). These shadows are projected by the avatar and the objects on the ground. An effective size of the objects was empirically determined to be $\frac{2}{3}$ of the size of the avatar's hands.

3.4 Training modes

The following training modes were designed on the basis of repeated interviews with the domain experts, in which they discussed the state of the art therapeutic approaches. In order to allow

new users to learn how to use the system, we introduced a two-minute long interactive tutorial to make them feel more comfortable and learn the basics. Three different training modes are provided, one of which is fully customizable by the therapist. At the end of each mode, simple statistics are provided to the user and therapist. We describe briefly each mode below.

3.4.1 Tutorial Mode

The Tutorial Mode is accompanied both by visual text appearing on screen and by a voiceover giving instructions to the patient. The two main types of avatar, the patient's and the therapist's, are presented to the user with graphical animations, then the user is asked to walk to the center of CAVE2 inside a virtual red circle, to the position she will have to maintain during the training (Fig. 3 right). The position of the user in this case is computed again by using the Kinect v2. Simple tasks of reaching objects follow, with visual and audio feedback in order to involve the user and accustom her to the system. The final part of the tutorial involves using the Wand controller in order to enable and use the graphical user interface, as well as the speech recognition feature to activate specific options.



Figure 3: Left: The shadows cast on the ground help the user perceive depth in the 3D environment. Right: Before starting a new session, a red virtual circle helps the user locate and maintain the correct position.

3.4.2 Random Objects Mode

The Random Objects mode is meant to be a quick way to run the application every time in a slightly different way. This mode generates, one at a time, a given number of virtual objects in random positions, with a new random seed at each execution. Objects, as mentioned before, are placed only in allowable positions in order not to be occluded. Our algorithm provides a dynamic timeout for each object that also considers the current behavior of the user: if the first time the patient doesn't reach the object in a predefined amount of time, the object vanishes and the next one appears; the system will now allow more time to the user for reaching successive objects. The main purpose of the timeout is to avoid particularly difficult positions for a specific user and to limit the impact of possibly unfortunate randomly-generated positions.

3.4.3 Progressive Mode

The Progressive Mode is very similar to the previous one, but focuses on generating progressively further objects in a predefined direction. The aim of this mode is to start with simple movements and make them more difficult little by little, in order to accustom the patient or to better see her improvements over time.

3.4.4 Custom Mode

The Custom Mode is completely configurable by the therapist through a file present on the

master node which operates the CAVE2 environment. The position of each object in space and its related timeout can be simply set by inserting the respective value in the system. This approach allows unlimited types of training, which can be more specific for the situation of each patient.

3.5 Trajectory Mode and Distorted Mode

Our application offers two additional ways to customize the user experience: a trajectory visualization mode and a distorted mode, which can be both activated at any moment with the Wand controller or using voice control.

3.5.1 Trajectory visualization

This feature is intended to improve the movements of the user. The feature employs visible trails which generate from the hands of the avatar and decay after a predefined amount of time (Fig. 4 left). This allows both the patient to be more mindful of her own movements and the therapist to notice significant details during the training. The trails are accompanied by a visualization of the optimal trajectory to reach a target: when a new virtual object appears, a ray is cast from the nearest hand of the avatar to that object, indicating the shortest path which could be followed to reach it. By visually comparing the white trails and the red trajectory, it very easy to identify qualitatively the movement error.



Figure 4: Left: When Trajectory Mode is enabled, the movements of the avatar generate white trails disappearing after a certain amount of time. At the same time, a red ray highlights the optimal trajectory to reach the current target. Right: In this example, the Distorted Mode maps the movement of the right part of the body of the patient to the left part of the avatar.

3.5.2 Distorted Reality

The current implementation of the *Distorted Reality* mode inverts the movements of the left and right parts of the body, so that if the user moves her left arm, the right arm of the avatar will move instead (Fig. 4 right). The main purpose of adding this feature was to study how patients respond when they are given an unusual situation or a new challenge that requires more concentration with respect to the normal task. The movement, in this case, requires a completely new approach from the patient, who typically finds it initially totally counter intuitive. A possible interesting application made possible with this feature would involve people who have limb amputations: for example, even if a patient does not have her left arm anymore, sometimes she may experience phantom pain – feeling pain to that limb as it was still there. Moving the right limb and seeing the left one moving instead in virtual reality could help the patient relieve that pain.

3.6 User Interaction

User interaction in the case of patients with upper-limb disabilities is a fundamental concern when designing the application.

Originally, the therapist was meant to activate the system and assist the patient during the whole session. The therapist is represented by her own avatar in the virtual scene, different

from the one of the patient. While the application is started from the master node of CAVE2, the Wand controller can be used for opening a menu panel. The panel is implemented as a virtual 3D user interface. This menu appears in the scene to the left of the avatar. The menu allows the user and the therapist to select which training mode to start (or restart), or to exit the application. Selections can be made by using specific buttons on the controller or by moving the avatar so that it touches with its hand the virtual interface.

However, we found this approach inefficient, as the CAVE2 environment renders the 3D perspective of the patient, and not the one of the therapist. Nevertheless, if we want to allow the patient to modify the training at run time, we have to consider that it is not always feasible for her to use a manual controller. To alleviate this problem, we added speech recognition with predefined words, that allows users to perform the functions that could have already been enabled with the Wand (Fig. 5).

For example, a patient can use her voice to change the training mode, restart it, close a menu, show the help panel, visualize the optimal trajectory or activate the distorted mode. Overall, since the different types of interaction do not interfere with each other, we let the user decide which interaction type she prefers. Animations, sound and visual effects are used in order to involve the patient and provide feedback (Fig. 6).



Figure 5: The user can use the Wand controller or his voice to interact with the 3D user interface. In this case, the tutorial helps the patient select the desired training mode.

3.7 Data Gathering

During each training session, the system records the position and rotation in space of each joint with a frequency of 30Hz. Also, additional information about the objects and the time in which they were reached are added to a log file, stored on the master node of the CAVE2. This information is then automatically sent to the laboratory, where data relative to a specific patient can be further analyzed. A different log for each mode is stored, even if incomplete, with basic metadata characteristic to the patient. The time spent by the patient in reaching a single target is saved and used as a performance measure, but other types of error measurements can be inferred and used from the recorded data (e.g., the relative angle to the optimal trajectory).

4 User Study

In order to test the efficiency of our application, we have conducted a pilot user study on 11 users without disabilities who acted as patients. They were given the same 5 tasks to complete in two separate trials, spaced one week apart. Each time they were asked to answer a questionnaire. All the testing was supervised by a senior domain expert.

The tasks proposed involved no external assistance and were as follows: 1) completing the *Tutorial Mode*, 2) enabling *Distorted Mode*, 3)



Figure 6: Animations, sound and visual effects are used in order to involve the patient and provide feedback.

completing a full training session in *Random Objects Mode*, 4) starting, restarting, changing and exiting a training session, 5) completing *Custom Mode* with *Movement Visualization* enabled, 6) using the voice to activate functions and the user interface. For each task we measured the task success, the time-on-task, and the number of errors.

The results show that all users completed each task successfully. During the first trail of tests, 4 users out of 11 sometimes wasted time in identifying the right buttons on the *Wand* controller or had to repeat the vocal commands for activating certain functions, especially because they did not pronounce the words loud enough. The number of errors for each task never exceeded 1, except in the case of voice recognition. The average time per tasks were as follows: 1) 208s 2) 63s 3) 46s 4) 15s 5) 19s 6) 23s.

During the second trial, we recorded only 2 overall errors. Furthermore, the average time-on-task was about 1/2 of the first trial one, very close to the one obtained by the developers themselves. This demonstrates that our avatar interface, despite relying on technology typically unfamiliar to the user, has high learnability and accustoms users with very little effort. The average time per tasks were as follows: 1) 66s 2) 48s 3) 25s 4) 6s 5) 14s 6) 17s.

The questionnaire indicates that the user experience was really good. 7 users out of 11 would not make any change to the current implementations. The *Trajectory Visualization* feature was considered to be very useful, as well

as the other features aimed at helping the depth-perception of objects in virtual reality. The *Distorted Reality* feature was surprising at first to some users. However, this mode proved to be learnable, after having collected just a few objects.

10 testers out of 11 believed that the scene was not distracting for the purposes of the task and that the non-photorealistic shape of the avatar did not influence the experience in any way. Two testers stated they would have preferred even bigger objects and UI elements, but the remaining felt comfortable with the graphical user interface, which was preferred to the use of voice in 8 cases out of 11. On top of these considerations, all the users judged the experience very fun and challenging.

The domain expert evaluated in detail the project, with very positive feedback, and intends to start testing the application on real patients as soon as possible, in order to conduct more specific case studies.

5 Discussion and Future Work

As far as tracking accuracy, Kinect v2 data can be noisy and should be filtered. Even light conditions and clothes worn by patients seemed to make positions less accurate. However, we can leverage the 10 infrared camera tracking system available in CAVE2 in order to have more precise data, possibly unaffected by other environmental elements thanks to the use of specific wearable trackers.

The existing interface can be further improved by the use of novel means of interaction in virtual reality, to facilitate interaction among multiple users, or between the patient and her therapist. For example, we are developing a special mode in which the therapist can move at run time the objects just by using her hands in the virtual space. A possible area of interest would be using our environment to test the possibilities of virtual rehabilitation on patients with amputations.

6 Conclusions

RehabJim offers a novel approach to the virtual rehabilitation of upper limbs, using a third per-

son perspective. The proposed framework leverages the advantages of Serious games in virtual reality environments and provides the possibility for therapists to interact directly with patients during a training session. According to the domain experts, this framework will be used to compare results of the third person perspective with a first person perspective, to better understand the merits of each approach.

References

- [1] Rehabilitation institute of chicago. <http://www.ric.org/about/>. Accessed: 2016-02-10.
- [2] Albert A Rizzo and Gerard Jounghyun Kim. A swot analysis of the field of virtual reality rehabilitation and therapy. *Presence*, 14(2):119–146, 2005.
- [3] JH Crosbie, Sheila Lennon, MDJ McNeill, and SM McDonough. Virtual reality in the rehabilitation of the upper limb after stroke: the users perspective. *Cyberpsychology & Behavior*, 9(2):137–141, 2006.
- [4] Paula Rego, Pedro Miguel Moreira, and Luis Paulo Reis. Serious games for rehabilitation: A survey and a classification towards a taxonomy. In *Information Systems and Technologies (CISTI), 2010 5th Iberian Conference on*, pages 1–6. IEEE, 2010.
- [5] Minhua Ma and Kamal Bechkoum. Serious games for movement therapy after stroke. In *Systems, Man and Cybernetics, 2008. SMC 2008. IEEE International Conference on*, pages 1872–1877. IEEE, 2008.
- [6] A Petracca, M Carrieri, D Avola, S Basso Moro, S Brigadoi, S Lancia, M Spezialetti, M Ferrari, V Quaresma, and G Placuir. A virtual ball task driven by forearm movements for neuro-rehabilitation. In *Virtual Rehabilitation Proceedings (ICVR), 2015 International Conference on*, pages 162–163. IEEE, 2015.
- [7] S Hoermann, L Santos, N Morkisch, K Jettkowski, M Sillis, NJ Cutfield, H Schmidt,

- L Hale, J Kruger, H Regenbrecht, et al. Computerized mirror therapy with augmented reflection technology for stroke rehabilitation: A feasibility study in a rehabilitation center. In *Virtual Rehabilitation Proceedings (ICVR), 2015 International Conference on*, pages 199–206. IEEE, 2015.
- [8] James William Burke, MDJ McNeill, Darryl K Charles, Philip J Morrow, Jacqui H Crosbie, and Suzanne M McDonough. Optimising engagement for stroke rehabilitation using serious games. *The Visual Computer*, 25(12):1085–1099, 2009.
- [9] Nick Mumford, Jonathan Duckworth, Patrick R Thomas, David Shum, Gavin Williams, and Peter H Wilson. Upper limb virtual rehabilitation for traumatic brain injury: Initial evaluation of the elements system. *Brain injury*, 24(5):780–791, 2010.
- [10] Nick Mumford, Jonathan Duckworth, Patrick R Thomas, David Shum, Gavin Williams, and Peter H Wilson. Upper-limb virtual rehabilitation for traumatic brain injury: A preliminary within-group evaluation of the elements system. *Brain injury*, 26(2):166–176, 2012.
- [11] Andrea Turolla, Mauro Dam, Laura Ventura, Paolo Tonin, Michela Agostini, Carla Zucconi, Pawel Kiper, Annachiara Cagnin, and Lamberto Piron. Virtual reality for the rehabilitation of the upper limb motor function after stroke: a prospective controlled trial. *Journal of neuroengineering and rehabilitation*, 10(1):1, 2013.
- [12] Lamberto Piron, Andrea Turolla, Michela Agostini, Carla Zucconi, Feliciano Cortese, Mauro Zampolini, Mara Zannini, Mauro Dam, Laura Ventura, Michela Battauz, et al. Exercises for paretic upper limb after stroke: a combined virtual-reality and telemedicine approach. *Journal of Rehabilitation Medicine*, 41(12):1016–1020, 2009.
- [13] Jigna Patel, Gerard Fluet, Alma Merians, Qinyin Qiu, Matthew Yarossi, Sergei Adamovich, Eugene Tunik, and Supriya Massood. Virtual reality-augmented rehabilitation in the acute phase post-stroke for individuals with flaccid upper extremities: A feasibility study. In *Virtual Rehabilitation Proceedings (ICVR), 2015 International Conference on*, pages 215–223. IEEE, 2015.
- [14] Alessandro Febretti, Arthur Nishimoto, Terrance Thigpen, Jonas Talandis, Lance Long, JD Pirtle, Tom Peterka, Alan Verlo, Maxine Brown, Dana Plepys, et al. Cave2: a hybrid reality environment for immersive simulation and information analysis. In *IS&T/SPIE Electronic Imaging*, pages 864903–864903. International Society for Optics and Photonics, 2013.
- [15] Mónica S Cameirão, I Badia S Bermúdez, E Duarte Oller, and Paul F Verschure. The rehabilitation gaming system: a review. *Stud Health Technol Inform*, 145(6), 2009.
- [16] Judith E Deutsch, Jennifer Rothman, Bryan Barker, Andrew Grando, and Harish Damodaran. The effect of video game interaction on walking intensity: Preliminary study of young, older adults and persons post-stroke. In *Virtual Rehabilitation Proceedings (ICVR), 2015 International Conference on*, pages 24–29. IEEE, 2015.