DualPanto: A Haptic Device that Enables Blind Users to Continuously Interact with Virtual Worlds

Oliver Schneider*, Jotaro Shigeyama, Robert Kovacs, Thijs Jan Roumen, Sebastian Marwecki, Nico Boeckhoff, Daniel Amadeus Gloeckner, Jonas Bounama, Patrick Baudisch
Hasso Plattner Institute
University of Potsdam, Germany
{firstname.lastname}@hpi.uni-potsdam.de

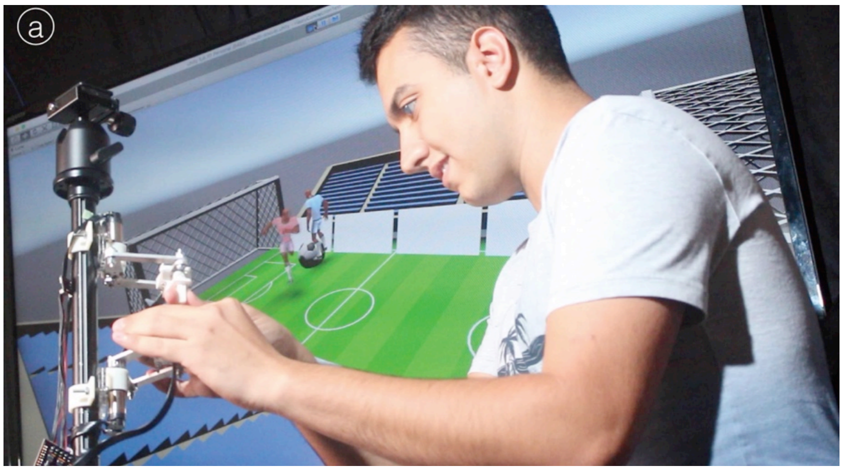

Figure 1: (a) This blind user is playing a real-time soccer game. (b) The DualPanto bimanual haptic device enables this. The user moves the *me* handle to move their avatar around the virtual world. Force feedback allows users to feel virtual walls, such as the edge of the field, and prevents users from pushing through them. The *it* handle moves by itself, enabling users to feel where the ball or opponent is located. Users aim by turning the haptic knob on the *me* handle; a foot pedal shoots the ball.

ABSTRACT

We present a new haptic device that enables blind users to continuously track the absolute position of moving objects in spatial virtual environments, as is the case in sports or shooter games. Users interact with DualPanto by operating the *me handle* with one hand and by holding on to the *it handle* with the other hand. Each handle is connected to a pantograph haptic input/output device. The key feature is that the two handles are *spatially registered* with respect to each other. When guiding their avatar through a virtual world using the *me* handle, spatial registration enables users to track moving objects by having the device guide the output hand. This allows blind players of a 1-on-1 soccer game to race for the ball or evade an opponent; it allows blind players of a shooter game to aim at an opponent and dodge shots. In our user study, blind participants reported very high enjoyment when using the device to play (6.5/7).

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

UIST '18, October 14–17, 2018, Berlin, Germany
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5948-1/18/10…$15.00
https://doi.org/10.1145/3242587.3242604

Author Keywords

Haptics; force-feedback; accessibility; blind; visually impaired; gaming.

ACM Classification Keywords

• **Human-centered computing~Haptic devices** • *Human-centered computing~Accessibility technologies*

INTRODUCTION

Today, blind users[1] operate a wide range of computer applications, from dynamic text documents [14] to games [22]. Spatial applications are also available through novel interfaces like tactile maps [29], mobile way finding devices [37], and graphics displayed through touch [10, 28].

However, there are no systems yet that would allow blind users to continuously track the absolute position of moving objects in virtual worlds, as would be necessary to play shooter or sports games. This class of applications is particularly challenging because it requires users to have a notion of a virtual space *and* to be able to track *moving* objects in the represented world, such as the player's opponents or a soccer ball.

---

[1] In this paper, we use "blind user" to mean our target user: someone unable to see the output of a computer screen.

*Oliver Schneider is now with the University of Waterloo.

Spatial environments without motion are well studied. Raised lines created using swell paper [43] or 3D printers [40] and Braille arrays [14] allow creation of dot and line drawings. The same strategy applies to spatial documents displayed using a force-feedback system, which users explore by scanning with an end effector, e.g., a stylus [28].

Unfortunately, blind users cannot perceive such spatial displays all at once, but must "scan" the space sequentially by running their fingers across it [17]. If anything changes or moves in such a display, blind users will not become aware of the change until they reach this object during the next scan. This makes it impossible to track moving objects, such as opponents or soccer balls, and essentially limits interaction to turn-taking.

To inform users about a moving object, researchers proposed adding a vibrotactile actuator that indicates the distance to a ball (e.g., [32]) or a stylus that points towards a target [33]. Similarly, audio pitch and volume can provide coarse information about an object's distance [1]. But, how can a device indicate the *location* of a moving object?

In this paper, we present a new approach that builds on *motion guidance* [48]. We combine two pantograph force-feedback devices in a spatially registered arrangement. The resulting device, *DualPanto*, enables users to continuously experience the spatial relationship between the user's avatar and other objects in the virtual world.

DUALPANTO

DualPanto is a haptic device that enables blind users to track moving objects while acting in a virtual world. Figure 1a shows a blind user using *DualPanto* to play a real-time soccer game.

As shown in Figure 1b, the device features two handles. Users interact with DualPanto by actively moving the *me* handle with one hand and passively holding on to the *it* handle with the other. DualPanto applications generally use the *me* handle to represent the user's avatar in the virtual world and the *it* handle to represent some other moving entity, such as the opponent in a soccer game.

Figure 2 shows the mechanics of DualPanto. The *me* handle is attached to a haptic pantograph: two motors that drive a closed linkage assembly. The *me* handle itself is a rotary knob, actuated by a smaller motor. The *it* handle is an identical mechanism mounted upside down such that the two handles face each other. This setup allows each of the handles to move laterally in their horizontal plane and to rotate; each handle thus supports three degrees of freedom.

Each haptic pantograph device provides force feedback with the two motors driving the pantograph and the third motor driving the handle. Each pantograph device senses its current position and orientation with three encoders: one attached to each of the two pantograph motors and one attached to the handle motor.

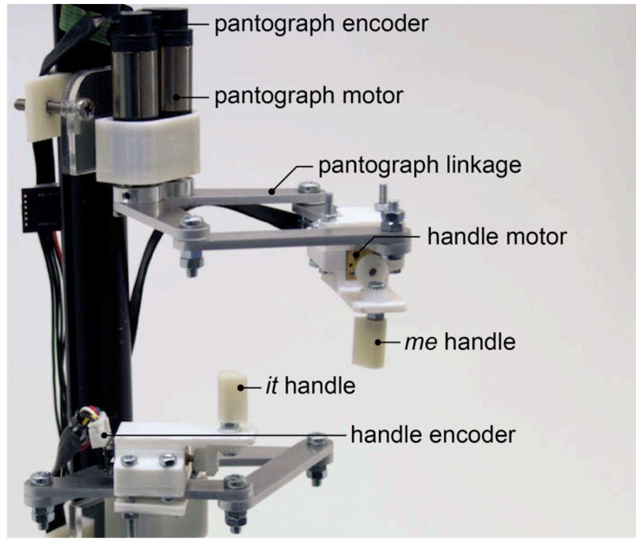

Figure 2: A DualPanto is built from two haptic pantographs, each of which has an actuated handle.

Interacting with DualPanto

The *me* and *it* handles continuously represent an avatar and another virtual entity, both of which can be moving. In Figures 3-7, we explain how it works in the context of our shooter game. In Figure 8 we explain how this interaction model extends to the soccer game shown in Figure 1.

***Me* handle** As shown in Figure 3, DualPanto's applications generally use the *me* handle to represent the user's avatar in the virtual world. (a) A user moves around and explores the virtual environment, such as a series of rooms in a shooter game, by moving the *me* handle. The device uses a direct 1:1 mapping, similar to the mapping of a touch screen; returning the handle to the same location in the physical world returns the avatar to same location in the virtual world. (b) The knob itself allows users to rotate their avatar. The pointed end of the handle represents the direction the avatar faces. (c) When the user pushes against a wall, the *me* handle resists by providing force feedback. (d) The *me* handle plays back haptic icons, for example, a short knockback when the player is hit by a projectile.

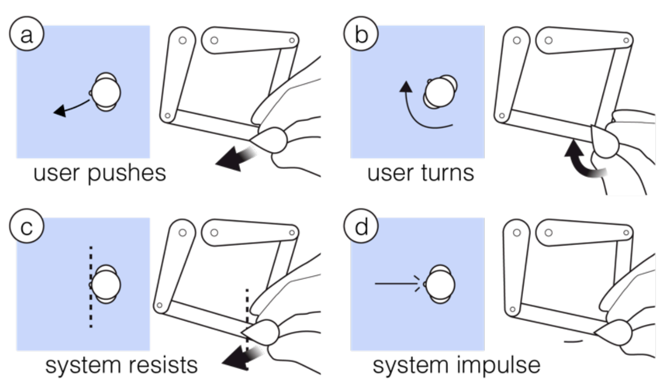

Figure 3: (a) The *me* handle allows users to move and (b) rotate their avatar. (c) The device renders walls and (d) impact using force feedback.

***It* handle** As shown in Figure 4, DualPanto applications use the *it handle* to render one selected moving object, such as the opponent in a first-person shooter. (a) If the object represented by the *it* handle moves, so does the handle. The handle is actuated *only* by the device, i.e., it will resist users trying to move it. By allowing the *it* handle to guide their hand, users can track the associated object, e.g., feel how the opponent moves and obtain a sense of where it is currently located with respect to the user's avatar. (b) At the same time, the actuated knob conveys to the user what direction *it* is facing. The *it* handle may also display haptic icons, such as an impulse when the opponent is shot.

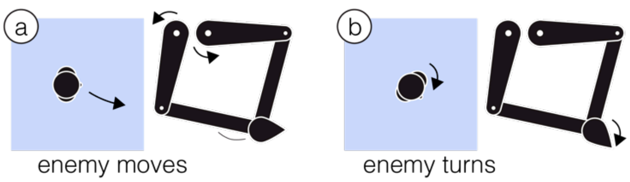

Figure 4: (a) Users track an object by allowing their hand to be guided by the *it* handle and (b) to be rotated by its knob.

Registration The key novelty behind DualPanto is that the *me* handle and the *it* handle are *spatially registered* with respect to each other, as illustrated by Figure 5. (a) For example, if the *me* handle is located directly *above* the *it* handle, the user knows that the user's avatar has collided with the opponent in the virtual world. In a shooter game, for example, this would typically indicate that the player is being attacked in hand-to-hand combat. (b) Spatial registration allows users to know where *it* is relative to *me*. In a shooter game this allows aiming and shooting at an opponent. (c) Similarly, feeling the position and orientation of the opponent allows users to evade the opponent's shots, e.g., by moving sideways while shooting (aka *strafing*).

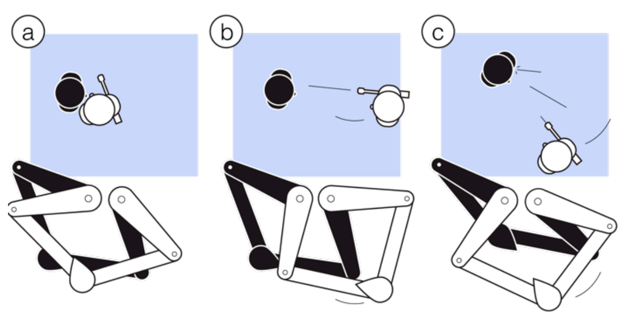

Figure 5: The *me* handle (white) and the *it* handle (black) are spatially registered. This allows users to interact in real-time. (a) When an enemy approaches the player, the user can (b) dodge while starting to take aim. (c) After aiming at the enemy, the player shoots, e.g., with a button or foot pedal.

Hiding and revealing *it* As shown in Figure 6, *it* may not be visible to the user. (a) When there is no visible *it*, e.g., there is no enemy, the *it* handle is relaxed: the motors driving this handle are turned off, allowing the user to back-drive the handle. (b) When *it* enters the user's line of sight, such as when the user enters a new room in a shooter game, the motors engage and snap the *it* handle to its position, revealing an *it* to track.

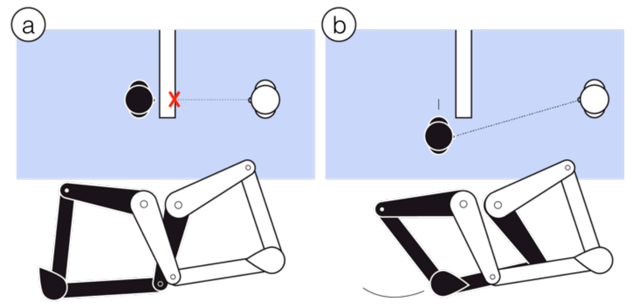

Figure 6: (a) When *it* is not visible to the user, the *it* handle is relaxed. (b) When *it* enters the user's line of sight, the handle snaps to *its* position.

Multiplexing *it* When the virtual scene contains multiple relevant objects, the *it* handle multiplexes between them. Figure 7 shows an example for a shooter game. (a) Here *it* may always represent the *closest* opponent. If the opponent is defeated, (b) the *it* handle automatically snaps to the next closest opponent.

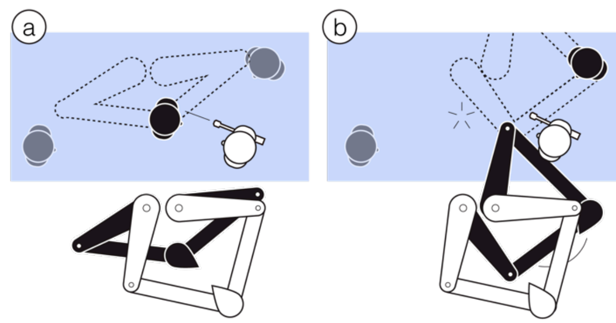

Figure 7: Multiplexing *it*. (a) The closest enemy is displayed on the *it* handle so that the user can take aim and shoot. If this opponent goes down, (b) the *it* handle snaps to the next closest opponent.

Applications

Conceptually, a wide range of applications for blind users can benefit from DualPanto, in particular any application that involves a spatial relationship between two or more entities, such as drawing curves with digital tape sketching [4] or a spatial sense-making application like a dynamic home finder [46] that can guide users to available homes and display routes to the user's workplace. However, DualPanto's ability to render objects moving in real time makes it particularly suited for applications that represent a virtual world, such as real-time spatial games. To illustrate this, we implemented two example games.

Shooter is a first-person survival shooter game. In this game (based on [42]) the player must fight a horde of zombie bunnies by shooting at each zombie as it approaches while avoiding their touch.

1-on-1 Soccer is the two-player game we implemented. Figure 8 shows an example interaction. (a) The opponent has the ball, so the player pays attention to which way they are looking. (b) The opponent shoots the ball; the *it* handle now follows the soccer ball as it moves. (c) Knowing where the ball is going, the player runs towards the ball and intercepts it. The player catches the ball and the *it* handle snaps to the opponent. (d) The player notices an opening and scores a goal by rebounding off the wall. (In our implementation, the goals span the width of the field.) Figure 1 shows a blind player performing this maneuver.

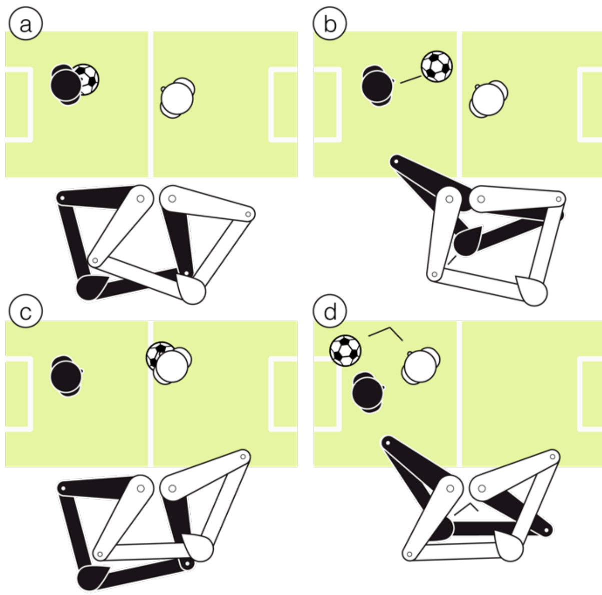

Figure 8: Walkthrough of our soccer game.

Audio feedback

We complement the haptic information from DualPanto with audio cues to provide additional information to the players during the game.

In the soccer game, our system provides audio cues for discrete changes in the game state: a user picking up the ball or shooting it, scoring a goal, and the ball reset to the center field. When a user holds the ball, "one" or "two" is spoken repeatedly to keep users aware of who has the ball.

In the first-person shooter, our system provides audio cues for footsteps for both player and zombies, gunshots, hits on a zombie (reinforced with haptic impulses on the *it* handle), and zombie attacks on the player (reinforced with haptic impulses on the *me* handle).

CONTRIBUTIONS & LIMITATIONS

Our main contribution is a new haptic device that enables blind users to interact with spatial virtual environments that contain moving objects without having to scan the environment. DualPanto accomplishes this by combining two paradigms of haptic feedback: an actively explored virtual environment for player-centric information, and a passively displayed stream of spatial information for tracked objects.

For this to work, the two handles must be spatially registrated. As we demonstrated in our walkthrough, spatial registration of the two handles allows for various types of interaction including spatial navigation, locating and tracking of virtual objects, dodging, hiding, and aiming.

Our device is limited in that it can only render two-dimensional worlds, and that tracking multiple objects requires it to time multiplex.

RELATED WORK

Our work builds on previous work in haptic environments, haptic display of streamed spatial data, and co-registered environments, especially for blind and visually impaired users.

Haptic feedback for exploring virtual environments

Haptic displays for blind and visually impaired users have traditionally focused on four main areas: Braille displays, tactile or haptic graphics, guidance and wayfinding, and sensory substitution [30]. Of these, Braille displays and tactile and haptic graphics are able to render virtual environments.

Virtual environments can be rendered for the tactile sense, i.e., sensed with the skin so that the point of connection is the user's finger and hands on a tactile display. Tactile graphics [10] are two-dimensional tactile renditions of information. Common static versions include raised graphics on Swell paper [43], buckled lines [36], or thermoform shapes produced by vacuum [10].

Dynamic versions have been rendered with haptic devices like skin-stretch actuators mounted on a pantograph mechanism [31], large Braille pin displays [14], or even larger 3-D printed lines [40]. Surface haptics involves tactile rending on surfaces, e.g., with programmable friction [3,47].

Virtual environments rendered with force feedback involve one or more end effector(s). For blind users, spatial data can be represented by means of haptification [28], for example, maps [27], or force-feedback line graphs [13], where their hands are guided along a line as they move the handle.

Gaming environments for blind and visually impaired users are dominated by audio games [2,21]. Spatial environments rely on mappings of location to pitch or volume [1] or spatial audio, e.g., using left/right panning [25] or more sophisticated binaural audio rendering [26]. Advanced sonification engines like AcouMotion can connect feedback to sensed body position, e.g., in a badminton game [24]. Voice interaction can involve spatial or postural information, for example, with eyes-free yoga [38]. Vibrating tactile cues can provide additional information, for example, timing in a tennis game [22], or location of non-moving objects (e.g., bowling pins) by scanning with a Wiimote in a bowling game [23]. Moving objects have been approximated with vibration, e.g., with a vibrating pong paddle [35].

The tightest coupling between controlling an avatar and tracking a moving object is with a single haptic stylus controller for Second Life [33]. The user provides relative input (e.g., as a joystick) while the stylus rotates to point towards a target technique; vibration intensity indicates distance.

No previous system has enabled absolute tracking of the users' avatar and a target, or tracking of a moving target's orientation, both of which are supported by DualPanto.

Streamed spatial information

An alternative to explored haptic environments, haptic systems can stream spatial information to users as output only. For example, work on haptic icons [20] found that space is a major design parameter [12]. Grids of vibrotactile actuators can produce the motion of phantom sensations and apparent motion [16] across actuators; this can be streamed in real-time [41]. Skin-drag displays [15] directly move a tactor across the skin for more salient motion.

Proprioception can also stream spatial information for users using motion guidance. Originally a technique for teaching motor skills like golf swings [48], motion guidance can also be display information. Gesture output [39] moves a user's fingers to display gestures. Muscle Plotter [19] uses electrical muscle stimulation to actuate users' arms to plot graphs. Lopes et al. also envisioned a proprioceptive hand-slap game, where one hand is actuated by the user and the other by the system using electrical muscle stimulation [18].

Motion guidance has been used to help blind users navigate with an actuated cane [5]. More recently, the McSig system helped teach handwriting to blind students through trajectory replay and real-time demonstration by a teacher [32].

Unlike streamed spatial display, DualPanto enables users to act within the same environment as the display.

Registered input and output

DualPanto uses a single registered workspace for locating both hands in the same space. Displays like LucidTouch [45] have done this for a visual touchscreen device having interaction both in front of and behind a mobile device. Bimanual haptic interfaces often allow both hands to act in a virtual environment, such as surgery training simulators [44]. However, each hand in these environments represents exactly the user's hands or implements. With DualPanto, in contrast, each hand corresponds to a different virtual object, and the user only acts in the virtual environment with the *me* hand.

IMPLEMENTATION

To assist readers in replicating our system, we discuss pantograph construction, handle design, software architecture, and haptic rendering. We developed the concept of *me* and *it* through brainstorming and rapid prototyping, and were informed by visually impaired collaborators during piloting. The cost of each DualPanto prototype was approximately 150USD; a production version could cost less.

Pantographs

DualPanto implements the haptic pantograph design [6,7]. We chose this design because it is planar and therefore appropriate for a single registered workspace without collisions; the two pantographs operate independently, unlike constructions where two pantographs control a single end-effector [8]. We based our implementation on the open-source Haply platform [11].

The two pantographs are mounted individually onto a tripod using a 3D-printed bracket. We typically weight the tripod's legs to improve stability. The linkages are made of 3 mm laser-cut aluminum, which is rigid enough to avoid deflection. This maintains the vertical distance between the two handles.

As illustrated in Figure 2, each of DualPanto's pantographs contains two DC motors (*Portescap*, 12V) with attached 27:1 gearheads and optical encoders. We chose the motors and gear ratio carefully to be able to overpower user's hand when rendering an opponent's position or a wall, but still to allow users to back-drive the motors in the *me* handle.

In order to achieve a reasonable resolution, we minimized the size of the handle and reduced slack by choosing motors with gearboxes that have little backlash and connecting linkages with double nuts. We used Teflon washers to cover the overlapping area between linkages.

Handles

The *me* and *it* handles are mounted to the arms of the pantographs by a 3D-printed bracket. The *me* handle contains a small motor (*Pololu*, 6V) with a gear ratio of 10:1, which is easily back-drivable for user input. The *it* handle has a higher gear ratio of 75:1 to provide enough force for system output.

To represent direction, we mounted a 3D-printed, asymmetric tip onto each of DualPanto's handles. After experimenting with several designs (Figure 9a) the "flattened teardrop" design (8x10x15 mm) performed best, in that holding this design between index finger and thumb clearly conveys its orientation any time (Figure 9b).

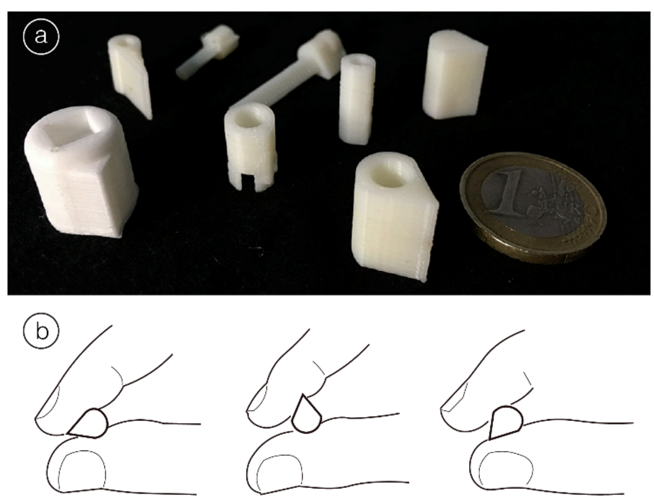

Figure 9: (a) After experimenting with various 3D printed shapes and sizes, (b) the "flattened teardrop" design performed best, because its orientation can be felt any time.

By arranging the two pantographs so that their handles point towards each other, we were able to get the two handles to be so close as to enable the two hands to touch during operation. This provides users with additional tactile cues that complement proprioception and improve the user's sense of registration.

Software architecture

DualPanto interfaces with the motors and encoders using an Arduino Due board and L298N motor driver. It uses native USB serial for high-speed communication to a PC (Macbook Pro). This enables us to run the haptic rendering on the computer. Our example games are implemented in Unity using the Java-based Haply framework [11] ported and extended into C#.

Haptic rendering

DualPanto employs four established haptic rendering techniques: (1) virtual coupling, (2) god-object constraints, (3) PD control, and (4) force impulses.

Virtual coupling [9] between the player's character and the end-effector provides resistance to the user when they try to move the *me* handle faster than the character can move.

DualPanto's end-effectors are implemented as custom Game Objects in Unity and coupled with god-object rendering [49] to simulate constraints, e.g., walls. These end-effectors are also coupled to in-game characters in our applications.

To output the correct amount of force we drive DualPanto's motors using traditional PD control loops. A second PD controller drives the rotation of the handle's knobs.

For both the *me* and *it* handles, fixed-value impulses provide feedback for collision events, such as picking up or shooting a soccer ball, being hit by a zombie, or a zombie being hit by a shot.

USER STUDY

We conducted a study to explore whether DualPanto enables blind users to interact with motion in spatial virtual environments. Participants played the 1-on-1 soccer game and reported on their experience. DualPanto is the first system to provide absolute position of targets in real-time; comparison with an existing system would lead to confounds. We thus chose an absolute evaluation to check whether blind users could play a real-time spatial game, and what that experience would be like.

Participants

We recruited six participants (ages 14-45, 1 woman) who had visual impairments. Participants' self-reported level of vision was between completely blind and 16% in one eye. Four participants were late-age blind. Participants' experience with video games varied from none at all to a maximum of 8h/week before blind (30min/month since). Having six participants let us analyze rich qualitative data.

Procedure

Participants played the 1-on-1 soccer game described earlier and shown in Figure 1a and Figure 8. We began by introducing and training each participant on the device and interaction model. First, we described the device and invited participants to freely explore. We introduced the *me* handle, first without power, and then with force-feedback so the participant could explore walls around the field. Next, we introduced the *it* handle by creating a ball on the field, tracked by the *it* handle. Participants first practiced moving to the ball to pick it up, and then shooting with a foot pedal. Finally, we introduced a second player for passing practice. At each stage, we waited until receiving confirmation that the participant was comfortable with introduced interaction model before moving on.

After training, participants played one or more matches against a trained player (one of the researchers) for 3-10 minutes and then provided their feedback in an interview. During the interview, we asked nine 1-7 Likert-like questions focused on experience, control, and spatial awareness to inform analysis and stimulate discussion.

One of the researchers took observations and transcribed quotations from video recordings. We clustered findings according to our a priori research question – does DualPanto enable users to interact with moving objects in spatial environments – and emergent themes from the data [34]. Questionnaire results are reported and informed analysis.

Results

Figure 10 shows the questionnaire results of the six participants, who reported a very high level of enjoyment (mean 6.50, standard error 0.22). All participants reported feeling in control (5.83 se 0.40) and when asked detailed questions about what elements they were able to track, generally agreed they had knowledge about their own location (5.17 se 0.48) and direction (5.67 se 0.71) and the ball's position (5.5 se 0.85), but less about the other player's location (4.66 se 0.95) and direction (5.17 se 1.05).

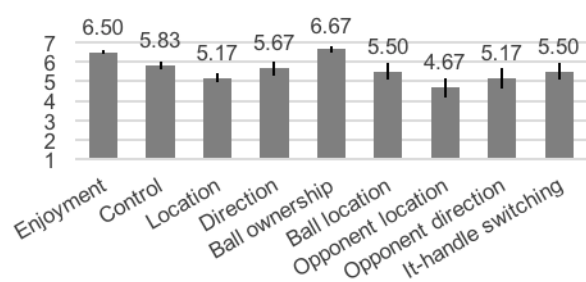

Figure 10: Questionnaire means and standard errors from our six participants, ranging from 1 (completely disagree) to 7 (completely agree).

Gameplay featured passes, blocks, and rebounds. Training lasted 8-15 minutes, and matches to 3 goals (the typical game played) lasted an average of 1m 48s (min 30s, max 3m 00s). Participants were positive about the device: *"it really gives you the opportunity to play, it pulls you in, like* *people say about computer games"* (P3), and reported having an enjoyable experience: *"it was over faster than I expected!"* (P4).

Qualitative results focus on spatial awareness, inclusive multiple experiences, and DualPanto as an audio-haptic gaming console.

DualPanto enabled spatial awareness of the game state

DualPanto enabled continuous interaction with moving objects. All six participants tracked their opponent and responded to their location, either to pass to them during practice (P2-6) or to shoot past them during gameplay (P1,2,4-6): *"yes I could definitely move the player to the ball and everything"* (P4). Participants could detect opponents' actions across the field: P3 remarked, *"that was on my goal"* when our trained player scored, and P5 noted when our trained player scored on his own net: *"own goal?"*

In addition to tracking the opponent, participants could orient themselves to the field. During training, P1 wanted to check his understanding: *"I would like to try an own goal."* He turned his character around to point directly to his net, and scored a goal in one shot, nodding as he heard the swoosh audio cue. After feeling where the two goals were, P3 was asked to go to the middle of the field; without hesitation, he moved his character to stand on the middle line.

Inclusive multiplayer experiences

Participants were surprised and excited about the prospect of multiplayer experiences: *"[multiplayer] is really significant, I have fun crushing people"* (P3), *"ah, cool that there are 2 players"* (P4). Participants also requested more multiplayer capabilities: *"it would be interesting to have more than 2 players"* (P3).

Multiplayer experiences evolved organically and included players with different levels of sight. After P3 beat our trained player 3-0 in 30 seconds, another sighted researcher exclaimed, *"okay, I'm playing!"* After defeating the second researcher in 47s, P3 asked the room with a laugh, *"next one?"* His sighted friend, who had accompanied P3 to the study, then received training and played with P3 for 3m48s. In a separate session, right before the match with our trained player, P5 called out with a grin, *"I will crush you, but you have to close your eyes!"*

DualPanto supported an audience and co-operative play. The game sessions with P5 and P6 was played in a gym during a blind soccer team practice. When P5 played, he gathered an audience who watched and listened to the game. P5 returned during P6's match and, standing over P5, started operating the DualPanto while P6 operated the foot pedal. P5 started telling P6 *"shoot!"* when he wanted P6 to press the pedal. Within a minute and a half of playing together, the two players were cooperating without needing to speak, other than shouts of excitement after goals.

DualPanto as an interactive audio-haptic gaming console

Audio provided context and engaged: *"I like…that you can instantly hear sounds, I really like it"* (P2). More sounds could improve the experience: *"it could be better if more sounds were added"* (P1), suggesting cheers and gasps from the crowd and sound effects from blind soccer.

Haptic feedback complemented the context provided by audio. Participants thought the combination could work for different games: *"you could explore levels with force feedback and show the next enemy with the it-handle, story is about audio anyways, so yeah I could imagine this working well for role play games"* (P3). Adding haptics improved interactivity: *"[another game] has different story lines that you can go into, but you always have to wait"* (P4).

Participants were receptive to a self-contained device, like an appliance or gaming console. Pick-up and play would be a valuable feature: *"for blind people it is nice when you have these things at home…because you don't need to adjust many settings"* (P2), *"I can just play for half an hour when I feel like it"* (P4). Some participants explicitly referred to it as a game console: P5 compared the interaction to a PlayStation, and P2-5 said that they would buy such a device: *"150 [euros]…that would be worth it, even if it only has a few minigames for now"* (P3).

As with existing gaming systems, such a console would require careful design for balance and flow. P1 both wanted slower gameplay ("*it was too fast to orientate"*) and more challenge (*"only shooting on the goal was too easy"),* suggesting rules to increase difficulty, like only counting corner shots. Tutorials and training would be a critical component: *"you need some practice until you can manage to do this, I noticed that I already got better"* (P4).

Additional feedback on ergonomics

Participants recommended ergonomic improvements to the hardware setup. P1 was uncomfortable using the device until we adjusted the arm rests and suggested additional support for the arms. Upon P4's request, we adjusted the handles to be further apart. P1 and P4 mentioned that the *it* handle could get in the way*: "it gets harder when the handles cross each other"* (P4).

Follow-up studies

Two follow-up studies gave us additional feedback.

Sighted participants

In our study, we noticed that DualPanto supported multiplay experiences among both sighted and blind users. In order to see how sighted participants respond to the device, we invited six additional sighted participants (ages 26-41, 2 women) to follow the same protocol with the soccer game. Feedback from sighted players was positive as well; participants reported a high level of enjoyment (mean 6.16, se 0.31) and high level of perceived control (mean 5.12, se 0.40). Thus, we found that sighted users enjoyed using DualPanto as well, and did not notice major differences in the gaming experience. This suggests that future research could also include sighted users' experiences.

Shooter game

We chose the soccer game for our user study because piloting suggested it would have broader appeal than a shooter, and because a multiplayer system facilitates training. As a follow-up, we invited an interested participant, P5, to experience the shooter game (Figure 11). P5 was able to play the shooter game and defeat multiple enemies, as shown in the accompanying video.

Figure 11: P5 playing our shooter game in a follow-up study.

DISCUSSION

DualPanto supported continuous interaction with moving objects in virtual worlds for both single and multi-user experiences. In our study, participants tracked themselves and opponents on the field, and interacted with them in real-time.

For DualPanto to work, audio and haptics needed to be designed in tandem. During development, we adopted a principle of rendering spatial content with haptics, and non-spatial content with audio; this approach was reinforced by user feedback. This suggests initial guidelines for other games: render story and framing with audio, continuous spatial content with haptics, and spatial events by using both modalities.

For DualPanto to render complex virtual worlds, it will need to offer enough resolution. Our current device ran simple games and could render multiple rooms, walls, or objects. Fitting more objects into the scene would require one of two solutions. 1) The workspace could be expanded, increasing absolute space. 2) The *resolution* could be increased by having higher-fidelity encoders and assembly to reduce slack. We expect these to have tradeoffs, and we plan to explore both in future work. Different devices might exist for different settings: a large desktop device may stay on a desk at home or work, but a smaller portable version may have enough resolution for simple tasks.

In the future, we envision DualPanto as a more general platform. Participants were already receptive to DualPanto as a platform for gaming, explaining the value of being able to pick up a console without adjusting anything and play for as short or long as they want. While we focused on game design as an ideal use case, DualPanto could support other applications.

For example, DualPanto could guide users through a web form, using the *me* handle to indicate the focused element and the *it* handle to display available form elements. DualPanto might work together with existing wayfinding applications: users could plan a hiking route at home with DualPanto, then export audio instructions to their smartphone. In an office setting, DualPanto could support information work like sketching, e.g., with digital tape [4]. We plan to explore these applications in future work.

CONCLUSIONS

We presented DualPanto, a haptic device that enables blind users to interact with spatial virtual environments that contain objects moving in real-time, as is the case in sports or shooter games. The key feature is that its two haptic in/output handles are *spatially registered* to each other, which enables blind players to navigate, track, dodge, and aim. In our user study, blind participants reported very high enjoyment when playing a soccer game (6.5/7).

Acknowledgements

This work was supported by the Hasso Plattner Institute and the Natural Sciences and Engineering Research Council of Canada (NSERC). We thank reviewers and participants for their valuable feedback.

REFERENCES

1. Audio Game Hub. http://www.audiogamehub.com
2. Audio Games. https://www.audiogames.net
3. Olivier Bau, Ivan Poupyrev, Ali Israr, and Chris Harrison. 2010. TeslaTouch: electrovibration for touch surfaces. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology (UIST '10)*. ACM, New York, NY, USA, 283-292. DOI:https://doi.org/10.1145/1866029.1866074
4. Ravin Balakrishnan, George Fitzmaurice, Gordon Kurtenbach, and William Buxton. 1999. Digital tape drawing. In *Proceedings of the 12th annual ACM symposium on User interface software and technology (UIST '99)*. ACM, New York, NY, USA, 161-169. DOI:http://dx.doi.org/10.1145/320719.322598
5. Johann Borenstein, and Iwan Ulrich. 1997. The GuideCane -a computerized travel aid for the active guidance of blind pedestrians. In *IEEE International Conference on Robotics and Automation (ICRA '97)*, Albuquerque, NM, USA, vol. 2, 1283-1288. DOI:https://doi.org/10.1109/ROBOT.1997.614314
6. Christophe Ramstein and Vincent Hayward. 1994. The pantograph: a large workspace haptic device for multimodal human computer interaction. In *Conference Companion on Human Factors in Computing Sys-*

tems (CHI '94). ACM, New York, NY, USA, 57-58. DOI: https://doi.org/10.1145/259963.260039

7. Gianni Campion, Qi Wang, and Vincent Hayward. 2005. The Pantograph Mk-II: A Haptic Instrument. In *Proc. IROS 2005, IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, pp. 723–728.
8. DiMaio, S. P., Salcudean, S. E., & Sirouspour, M. R. 2000. Haptic interaction within a planar environment. In *Proceedings of the 2000 ASME International Mechanical Engineering Congress & Exposition, Dynamic Systems and Control Division*. pp. 69-2.
9. J. Edward Colgate, Michael C. Stanley, and J. Michael Brown. 1995. Issues in the haptic display of tool use. In *International Conference on Intelligent Robots and Systems 95.'Human Robot Interaction and Cooperative Robots'(IROS '95)*. IEEE, vol. 3,140-145. DOI:https://doi.org/10.1109/IROS.1995.525875
10. Polly Edman. Tactile graphics. *American Foundation for the Blind,* 1992.
11. Colin Gallacher, Arash Mohtat, Steve Ding, and József Kövecses. 2016. Toward open-source portable haptic displays with visual-force-tactile feedback colocation. In *Haptics Symposium (HAPTICS)*. IEEE, 65-71. DOI:https://doi.org/10.1109/HAPTICS.2016.7463157
12. Eric Gunther and Sile O'Modhrain. 2003. Cutaneous grooves: composing for the sense of touch. *Journal of New Music Research* 32, no. 4, 369-381.
13. Mikro Horstmann, Cornelius Hagen, Alisdair King, Sijo Dijkstra, David Crombie, David Gareth Evans, George T. Ioannidis, Paul Blenkhorn, Otthein Herzog, and Christoph Schlieder. 2004. TeDUB: Automatic interpretation and presentation of technical diagrams for blind people. In *Proceedings of Conference and Workshop on Assistive Technologies for Vision and Hearing Impairment (CVHI '04)*. DOI:https://doi.org/10.1080/13614560512331326017
14. HyperBraille, http://www.hyperbraille.de/
15. Alexandra Ion, Edward Jay Wang, and Patrick Baudisch. 2015. Skin Drag Displays: Dragging a Physical Tactor across the User's Skin Produces a Stronger Tactile Stimulus than Vibrotactile. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 2501-2504. DOI: https://doi.org/10.1145/2702123.2702459
16. Ali Israr and Ivan Poupyrev. 2011. Tactile brush: drawing on skin with a tactile grid display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*, New York, NY, USA, 2019-2028. DOI: https://doi.org/10.1145/1978942.1979235
17. Susan J Lederman, Roberta L Klatzky. 1987. Hand movements: A window into haptic object recognition. In *Cognitive Psychology*, Science Direct, Volume 19, Issue 3, 342-368. DOI:https://doi.org/10.1016/0010-0285(87)90008-9
18. Pedro Lopes, Alexandra Ion, Willi Mueller, Daniel Hoffmann, Patrik Jonell, and Patrick Baudisch. 2015. Proprioceptive Interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 939-948. DOI:https://doi.org/10.1145/2702123.2702461
19. Pedro Lopes, Doăa Yüksel, François Guimbretière, and Patrick Baudisch. 2016. Muscle-plotter: An Interactive System based on Electrical Muscle Stimulation that Produces Spatial Output. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. ACM, New York, NY, USA, 207-217. DOI:https://doi.org/10.1145/2984511.2984530
20. Karon MacLean and Mario Enriquez. 2003. Perceptual design of haptic icons. In *Proceedings of EuroHaptics*. IEEE, 351-363. DOI:https://doi.org/10.1109/HAPTICS.2006.171
21. Daniel Miller, Aaron Parecki, and Sarah A. Douglas. 2007. Finger dance: a sound game for blind people. In *Proceedings of the 9th international ACM SIGACCESS conference on Computers and accessibility (Assets '07)*. ACM, New York, NY, USA, 253-254. DOI:http://dx.doi.org/10.1145/1296843.1296898
22. Tony Morelli, John Foley, Luis Columna, Lauren Lieberman, and Eelke Folmer. VI-Tennis: a vibrotactile/audio exergame for players who are visually impaired. 2010. In *Proceedings of the Fifth International Conference on the Foundations of Digital Games (FDG '10)*. ACM, New York, NY, USA, 147-154. DOI:http://dx.doi.org/10.1145/1822348.1822368
23. Tony Morelli, John Foley, and Eelke Folmer. 2010. Vi-bowling: a tactile spatial exergame for individuals with visual impairments. In *Proceedings of the 12th international ACM SIGACCESS conference on Computers and accessibility (ASSETS '10)*. ACM, New York, NY, USA, 179-186. DOI:http://dx.doi.org/10.1145/1878803.1878836
24. Hermann, T., Höner, O., & Ritter, H. (2005, May). AcouMotion–an interactive sonification system for acoustic motion control. In *International Gesture Workshop*. Springer, Berlin, Heidelberg. 312-323. DOI:https://doi.org/10.1007/11678816_35
25. Shades of Doom, http://www.gmagames.com/sod.html
26. Papa Sangre, http://www.papasangre.com/
27. Konstantinos Moustakas, Georgios Nikolakis, Konstantinos Kostopoulos, Dimitrios Tzovaras, and Michael G. Strintzis. 2007. Haptic rendering of visual data

for the visually impaired. *IEEE MultiMedia* 14, no. 1. DOI:https://doi.org/10.1109/MMUL.2007.10

28. Sabrina Paneels and J. C. Roberts. 2010. Review of Designs for Haptic Data Visualization. in *IEEE Transactions on Haptics*, vol. 3, no. 2, 119-137. DOI:https://doi.org/10.1109/TOH.2009.44
29. Peter Parente, and Gary Bishop. 2003. BATS: the blind audio tactile mapping system. In *Proceedings of the ACM Southeast regional conference (ACMSE '03)*. ACM, New York, NY, USA, 132-137.
30. Dianne TV Pawluk, Richard J. Adams, and Ryo Kitada. 2015. Designing haptic assistive technology for individuals who are blind or visually impaired. *IEEE transactions on haptics* 8. IEEE, no. 3: 258-278. DOI:https://doi.org/10.1109/TOH.2015.2471300
31. Grégory Petit, Aude Dufresne, Vincent Levesque, Vincent Hayward, and Nicole Trudeau. 2008. Refreshable tactile graphics applied to schoolbook illustrations for students with visual impairment. In *Proceedings of the 10th international ACM SIGACCESS conference on Computers and accessibility (ASSETS '08)*. ACM, New York, NY, USA, 89-96.DOI:https://doi.org/10.1145/1414471.1414489
32. Beryl Plimmer, Andrew Crossan, Stephen A. Brewster, and Rachel Blagojevic. 2008. Multimodal collaborative handwriting training for visually-impaired people. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 393-402. DOI:https://doi.org/10.1145/1357054.1357119
33. de Pascale M., Mulatto S., Prattichizzo D. (2008) Bringing Haptics to Second Life for Visually Impaired People. In *Ferre M. (eds) Haptics: Perception, Devices and Scenarios. EuroHaptics 2008*. Lecture Notes in Computer Science, vol 5024. Springer, Berlin, Heidelberg. DOI: https://doi.org/10.1007/978-3-540-69057-3_112
34. Gery W. Ryan and H. Russell Bernard. 2003. Techniques to Identify Themes. *Field Methods*, SAGE Journals, Vol 15, Issue 1, 85–109. DOI:https://doi.org/10.1177/1525822X02239569
35. Steven Strachan, Wiertlewski, M., Zophoniasson, H., and Anastassova, M. 2013. ViPong: Probabilistic haptic feedback for eyes-free interaction. In *IEEE World Haptics Conference (WHC).* IEEE, 193-198. DOI: https://doi.org/10.1109/WHC.2013.6548407
36. Raised Line Drawing Kits: http://www.maxiaids.com/raised-line-drawing-kit
37. David A. Ross, and Bruce B. Blasch. 2000. Wearable interfaces for orientation and wayfinding. In *Proceedings of the fourth international ACM conference on Assistive technologies (ASSETS '00)*. ACM, New York, NY, USA, 93-200. DOI:https://doi.org/10.1145/354324.354380
38. Kyle Rector, Cynthia L. Bennett, and Julie A. Kientz. 2013. Eyes-free yoga: an exergame using depth cameras for blind & low vision exercise. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '13)*. New York, NY, USA. DOI:http://dx.doi.org/10.1145/2513383.2513392
39. Anne Roudaut, Andreas Rau, Christoph Sterz, Max Plauth, Pedro Lopes, and Patrick Baudisch. 2013. Gesture output: eyes-free output using a force feedback touch surface. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2547-2556. DOI: https://doi.org/10.1145/2470654.2481352
40. Saiganesh Swaminathan, Thijs Roumen, Robert Kovacs, David Stangl, Stefanie Mueller, and Patrick Baudisch. 2016. Linespace: A Sensemaking Platform for the Blind. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 2175-2185. DOI: https://doi.org/10.1145/2858036.2858245
41. Oliver S. Schneider, Ali Israr, and Karon E. MacLean. 2015. Tactile Animation by Direct Manipulation of Grid Displays. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15)*. ACM, New York, NY, USA, 21-30. DOI: https://doi.org/10.1145/2807442.2807470
42. *Survival Shooter Tutorial* “Zombunny”, Unity Asset store. https://www.assetstore.unity3d.com/en/#!/content/40756
43. Swell-form-graphics: http://www.americanthermoform.com/product/swell-form-graphics-ii-machine/
44. Anthony Talvas, Maud Marchal, and Anatole Lécuyer. 2014. A survey on bimanual haptic interaction. *IEEE transactions on haptics* 7, IEEE, no. 3: 285-300.DOI: https://doi.org/10.1109/TOH.2014.2314456
45. Daniel Wigdor, Clifton Forlines, Patrick Baudisch, John Barnwell, and Chia Shen. 2007. Lucid touch: a see-through mobile device. In *Proceedings of the 20th annual ACM symposium on User interface software and technology (UIST '07)*. ACM, New York, NY, USA, 269-278. DOI: https://doi.org/10.1145/1294211.1294259
46. Christopher Williamson and Ben Shneiderman. 1992. The dynamic HomeFinder: evaluating dynamic queries in a real-estate information exploration system. In *Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval (SIGIR '92)*. ACM, New York,

NY, USA, 338-346. DOI: https://doi.org/10.1145/133160.133216

47. L. Winfield, J. Glassmire, J. E. Colgate and M. Peshkin. 2007. T-PaD: Tactile Pattern Display through Variable Friction Reduction, *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC '07)*. 421-426. DOI:https://doi.org/10.1109/WHC.2007.105
48. Yasuyoshi Yokokohji, Ralph L. Hollis, Takeo Kanade, Kazuyuki Henmi, and Tsuneo Yoshikawa. 1996. Toward machine mediated training of motor skills. Skill transfer from human to human via virtual environment. In *5th IEEE International Workshop on Robot and Human Communication*. IEEE, 32-37. DOI: https://doi.org/10.1109/ROMAN.1996.568646
49. C. B. Zilles, and J. K. Salisbury. 1995. A constraint-based god-object method for haptic display. In *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots*. IEEE, Vol. 3, 146–151. DOI:http://doi.org/10.1109/IROS.1995.525876