WAVI: Improving Motion Capture Calibration Using Haptic and Visual Feedback

Zlatko Franjcic[1,2*], **Paweł W. Woźniak**[2,4*], **Gabrielė Kasparavičiūtė**[3], **Morten Fjeld**[2]
[1]Qualisys AB & Chalmers University of Technology, Gothenburg, Sweden, zlatko.franjcic@qualisys.se
[2]t2i Lab, Chalmers Univeristy of Technology, Gothenburg, Sweden, zlatco,pawelw,fjeld@chalmers.se
[3]CSE, University of Gothenburg, Gothenburg, Sweden
[4]VIS, University of Stuttgart, Stuttgart, Germany, pawel.wozniak@vis.uni-stuttgart.de

ABSTRACT
Motion tracking systems are gaining popularity and have a number of applications in research, entertainment, and arts. These systems must be calibrated before use. This process requires extensive user effort to determine a 3D coordinate system with acceptable accuracy. Usually, this is achieved by rapidly manipulating a calibration device (e.g. a calibration wand) in a volume for a set amount of time. While this is a complex spatial input task, improving the user experience of calibration inspired little research. This paper presents the design, implementation, and evaluation of WAVI — a prototype device mounted on a calibration wand to jointly provide visual and tactile feedback during the calibration process. We conducted a user study that showed that the device significantly increases calibration quality without increasing user effort. Based on our experiences with WAVI, we present new insights for improving motion tracking calibration and complex spatial input.

Author Keywords
Calibration; Motion Tracking; Prototyping; Spatial Input

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION
Motion tracking systems enable their users to log the position of objects in space. These systems use images from multiple cameras to reconstruct the 3D positions of common points identified across the images. Users can define the volume to be tracked within and the objects to be tracked. The latter is most often achieved by placing reflective markers on the objects.

Marker-based motion tracking is becoming more accessible and its user group is widening. Filmmaking and gaming now depend on accurate and fast motion tracking to animate models and to create new visual environments. Many research

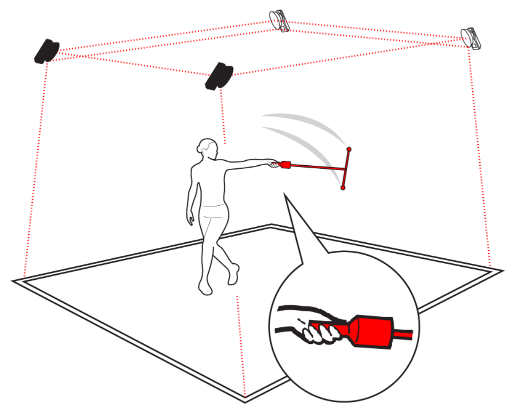

Figure 1: An overview of calibrating a motion capture system using WAVI. The user performs the wand dance procedure to calibrate the system while WAVI provides in-calibration feedback.

fields, such as biomechanics or cognitive psychology, depend on motion tracking to monitor human movement and understand more about the human body and mind. Sports scientists and professional athletes benefit from accurate analyses of the athletes' movements and look for possible improvements in posture. With a growing number of applications and new uses emerging, motion tracking systems need to become accessible to more users and offer a trouble-free experience.

In this work, we look specifically at the initial phase of using a motion tracking system — the calibration of large volumes. In order to establish a 3D coordinate system from 2D camera images, the system needs to acquire enough data to jump start the tracking algorithms and offer an acceptable accuracy. This is usually achieved by manipulating an object of known dimensions in the tracked volume (e.g. a wand), often with an additional object placed at a fixed location for reference. In order to attain the desired tracking quality, users must gather as many points as possible in their desired tracked volume. For large volumes, the user needs to physically walk into the volume holding the wand. In such a case, the user's body often obscures some of the cameras' views, which results in the need to dynamically walk within the tracked volume while moving the wand in all possible directions. The volume should be filled uniformly and as randomly as possible. This is sometimes called the wand dance procedure [28].

[*] Z. Franjcic and P. W. Woźniak are joint first authors of this work.

MobileHCI '16, September 06–09, 2016, Florence, Italy

ACM 978-1-4503-4408-1/16/09...$15.00
DOI: http://dx.doi.org/10.1145/2935334.2935374

While the movements are easy for experts, they are not so for a new user.

Here, we investigate whether the calibration process, specifically for large volumes, can be made more accessible to users and if calibration aids may be designed that will improve the calibration quality without an increase in user effort. We designed, built, and evaluated WAVI; a proof-of-concept augmentation for a standard calibration wand that provides visual and haptic feedback during the calibration process. While WAVI is built to work with a specific motion tracking system, it uses the features of the system that are shared with most commercial models. WAVI is fully integrated into the calibration wand so that the calibration procedure is not altered, as shown in Figure 1. This paper contributes the following: (1) the design and implementation of a device for in-calibration feedback for motion tracking, (2) an evaluation of the device through a user study with 20 participants, (3) insights from the design and evaluation process and (4) challenges for further research in feedback for motion tracking calibration. This paper is organised as follows. First, we present research related to motion tracking and 3D input. Next, we introduce the details of the design and implementation of WAVI. We then report on our user study and discuss its results. Finally, we point to interesting insights emerging from the study and explicate the requirements of a future system providing feedback during motion tracking calibration.

RELATED WORK

Marker-based motion capture has a number of application domains. Its growing availability and its importance as a tool in HCI (e.g. to understand pointing [34], build underwater companions [35], or to evaluate the ergonomics of touch surfaces [4]) are the main motivations for our work. There is little past research on the HCI aspects of motion tracking calibration. Consequently, our work builds on experiences of calibrating other kinds of devices. As the calibration is performed in 3D space, work on 3D input tasks, especially 3D drawing (due to similarities in the input task), inform our work. We also reviewed how devices similar to those used for calibration (wands) were used in past inquires. Finally, we reviewed how output during complex 3D tasks was designed in past research to look for effective ways to communicate calibration feedback to the user.

Calibration tasks

Significant advances in vision-based 3D reconstruction have been made in the past [22]. Improving calibration and reconstruction accuracy continue to be active areas of research [23, 32]. Novel calibration algorithms [30, 33], and replacing bulky and expensive 3D calibration objects with 1D wands have contributed to reducing manual calibration effort [28]. However, little work has focused on improving the procedure or user experience of the calibration. This is quite surprising as the random movements that are best generated by humans are still required to assure accuracy and reliable operation of a motion tracking system. One also needs to remember that the task is performed in a specific context of the measurement and the calibration procedure may need to be modified. For example, Fedotova and Pilipivs [7] built a special system designed to track bobsleighers in a training environment. While some attempts to replace humans with robots have been made [40], motion capture systems will still require user calibration in the foreseeable future. While a blocking calibration model can be applied (i.e. the system will not be operational unless a predefined accuracy is available [9]), this may cause frustration. It remains a challenge to determine the methods to calibrate systems efficiently and in a user-friendly way. Consequently, our work focuses on improving the quality and user experience of user calibration.

Several attempts at modifying calibration procedures are worth noting. Hoßbach et al. [16] used a set of spheres filled with water to make calibration more effective and user-friendly. Their approach, however, is strictly limited to motion tracking in Medical Resonance Imaging. Our work is also inspired by developments in improving the calibration processes for systems other than motion tracking. Lu and Huenerfauth [19] conducted an extensive inquiry into the calibration procedure of an interactive glove. Their work shows that a structured inquiry into how the calibration is organised and communicated to the user can result in significant improvements in user satisfaction. Pfeuffer et al. [26] improved the accuracy and usability of calibrating an eye tracker by redesigning user feedback and the procedure. Similarly, Villanueva et al. [38] proposed new mathematical models to reduce frustration and physical demand. Flatla et al. [8] used a gamification approach to create a design framework to disguise calibration tasks as games. In this way, calibration was found to be more enjoyable, but we were unable to apply it in our case as a gamified approach may have been inappropriate in professional contexts where motion capture systems are mostly used. Furthermore, the framework does not cover the kind of calibration task for which we were designing (i.e. filling a volume in a randomised manner). While all of the above work shows that inquiries into calibration procedures are meaningful and can lead to an enhanced understanding of user interaction, no past research investigated possible improvements to the calibration of motion tracking systems. WAVI goes beyond past work in calibration and focuses on a spatially complex task in the context of increasingly popular marker-based motion tracking systems. Our work is targeted at systems that use the wand dance calibration technique [2]. While we use a marker-based reconstruction system, the wand-based procedure can be used to calibrate marker-less optical systems as well [6]. Consequently, our work addresses a wider range of motion capture systems.

3D drawing

Given the characteristics of the wand dance procedure, we endeavored to identify input tasks that resembled the calibration procedure. Drawing (or painting) in 3D is of particular interest here, as the calibration task can be interpreted as painting the target tracking volume from the cameras point of view. While most work in this area focuses on sketching and curve editing (e.g. [5]), we are inspired by work that investigated painting larger virtual volumes. Schkolne et al. [31] showed that users can effectively use tangible tools to quickly create large volumetric shapes. Gregory et al. [10] investigated how three-dimensional meshes can be subject to user modification. They determined that force feedback can enhance volume perception for users and enable modifying and manipulating 3D shapes effectively. Similarly, McMains [20] showed that haptic feedback can be effectively used to modify 3D geometries. While we cannot directly apply the lesson learnt here in designing 3D drawing applications, we recognise the similarity between the wand dance calibration procedure and creating a 3D painting. Consequently, the aforementioned works inspire our design, especially in terms of feedback modality.

Devices for 3D interaction

While little research specifically addresses the calibration problem, input tasks based on manipulating objects in unrestricted 3D space have been explored in the past. AHNE [25] investigated how different objects can be placed and repositioned in space to produce a desired sound output. XWand [39] used an interactive wand with visual feedback to allow for richer interaction with rich environments. While all the previously developed systems mentioned above facilitate complex input in 3D environments and show how different forms of device feedback can be used, the work focuses on instantaneous pointing or interacting with objects. Interaction with volumes was limited to small entities such as in HoloDesk [15]. In contrast, we investigate interaction with large volumes and over an extended period of time.

Xiao and Balakirishnan [41] used a wand to integrate with a large screen, but they decided not to provide any feedback through the device. The opposite design choice was taken by Han et al. [13] who demonstrated that a 3D shape can be effectively perceived through a 2D haptic device. This concept was further explored in an extensive in-the-wild study by van der Linden et al. [36] where users navigated a pitch-black room with a handheld haptic device. Haptic feedback was also proven for the more precise task of 3D pointing by Grossman and Balakrishnan [11]. While WAVI addresses an interestingly different task, our research is influenced by past successes of applying haptic feedback in 3D input tasks. Our work investigates the applicability of haptic feedback in interacting with a large volume.

DESIGN

Inspired by [18], we share insights from our design process to create intermediate-level knowledge. Our work explores design possibilities for in-calibration feedback that could help users achieve better results when calibrating motion capture systems. As little past work in the area is available, we needed to conceptualise the design of WAVI by first analyzing the characteristics of the calibration task. We then proceeded to identify the possible feedback modalities and chose the ones to include in the prototype. Finally, we gave the device a physical form and integrated it with the calibration tools required by motion capture systems.

Calibration task

First, we took a closer look at what kind of input is required to optimally calibrate a motion tracking system. We analyzed the Qualisys Oqus system [29] as a product representative of a number of marker-based motion capture systems from different manufacturers. While other systems use different hardware, their features are very similar [24, 37]. As we did not intend to modify the calibration procedure, we looked at the existing tools and methods.

Figure 2: The two reference objects provided by the supplier of the motion capture system. The wand (left) and the L-frame (right).

The calibration requires two physical tools equipped with reflective markers to be placed in the calibrated space (Figure 2). An L-frame is placed at a fixed location to mark the capture volume's coordinate system origin. A T-shaped wand is to be manipulated by the user within the volume. After analyzing the user manual [29] and consulting with Qualisys tech support, we determined that the goal of the calibration task, from the motion capture system's point of view, was twofold. Firstly, the system should obtain as many points as possible during the calibration. Secondly, the user should place the wand in the entire volume they wish to be tracked. The system provides calibration quality feedback only at the end of the calibration process in the form of number of points collected, camera residuals and standard deviation of wand length. Advanced users of motion capture can intuitively judge if the numbers satisfy their requirements, while novice users need to determine that experimentally.

Next, we analyzed the user manual for instruction regarding the kind of movement needed for best calibration. There is little information about what movement is best, but rotating the wand with moderate speed is suggested. A technique for calibrating the volume near walls and floors is also provided. Two of the authors of this work spent significant time working with the system, practising calibration techniques and observing the produced quality measures to gain experience in calibration. We observed that better results were achieved with dynamic movement and that it was hard to enter all regions of the desired calibration volume. This enabled us to begin considering alternatives for in-calibration feedback.

Exploring the feedback design space

As the calibration already requires external devices, we decided for the feedback system to remain independent of additional outside infrastructure. That is, while we consider modifying existing external devices, we reject adding an additional device as this would further complicate the process which already requires carrying the calibration equipment in a dedicated case. We established not extending the calibration infrastructure as our key design principle. Instead, we considered augmenting one of the existing calibration tools. We also rejected any solutions that were not implementable within a predictable technical horizon in motion capture such as augmented reality glasses (as the initial position of the user is unknown in current systems when the system is not calibrated so proper visualisation for augmented reality is not possible). As the L-frame is bulky and stationary, our attention turned to the calibration wand. The user is in direct contact with the wand throughout the calibration process which makes the wand particularly suited for communicating feedback.

We then investigated what output modalities could be used for in-calibration feedback. Visual feedback on a screen was one possibility that would imply the lowest implementation costs. We decided to dismiss this design as on-screen feedback cannot be easily viewed by a user who is rotating their body and manipulating the wand. Checking the calibration status on a screen would require interrupting the task thus lowering calibration quality.

We also considered implementing visual feedback on camera-mounted displays similar to commercially available solutions [24, 37]. This kind of feedback relies on measuring 2D image coverage and counts of calibration points on a per-camera basis. While per-camera visual feedback indicating the calibration progress for the respective camera appears to be correct from a user perspective, the underlying concept of 3D motion capture strongly suggests [24, 29, 37] that the primary goal of calibration is to fill the target volume with samples of varying poses of the wand. The per-camera wand coverage and count give an incomplete substitute view of the calibration state, while the true goal is to thoroughly fill the target volume. Consequently, we dismissed per-camera feedback and instead opted for volume-based feedback for WAVI.

We further contemplated using audio feedback. For example, the system could produce differently pitched sounds depending on the calibration quality within the immediate vicinity of the user. We rejected this idea as it violated our principle of not adding additional devices, and speakers, 3D audio systems or headphones would add complexity to the calibration infrastructure. Another idea was to mount a speaker on the wand. This, however, would require a separate study as the speed of the wand movement (displacement and rotation) and the speaker's relative position to the floor (sound reflection) would have a direct, and possibly misleading, impact on the perceived sound signal. Furthermore, persistent audio feedback requires careful design as it is often considered annoying [14, 21].

Next, inspired by past work mentioned earlier in this paper, we directed our attention to haptic feedback. As the user is handling the wand throughout the calibration, it is possible to constantly communicate haptic feedback. Since vibration-based feedback has been preferred to pressure-based output in past work [25, 36], we decided to design for vibration feedback on the wand handle.

Finally, we considered visual feedback placed directly on the calibration wand. This modality has the advantage that one usually looks in the direction of the wand while calibrating the volume. The wand is carried before the user while moving in the volume based on the owner's manual and the authors experiences. This is a significant difference compared to feedback on a stationary screen. Furthermore, the wand is frequently rotated and consequently, the on-wand display cannot be complex and it must be meaningful in high-speed rotation. Based on these considerations, we chose to design low-fidelity visual feedback on the wand. We were then ready to assemble a prototype of WAVI.

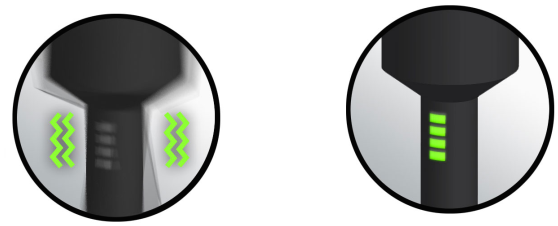

(a) Feedback for an entirely uncalibrated volume. WAVI vibrates with full power and no lights are displayed. (b) Feedback for a fully calibrated volume. No vibration is present and all lights are on.

Figure 3: Mapping of calibration levels to visual and tactile feedback in WAVI.

Prototype design

WAVI takes the form of a device that is permanently mounted on the handle of a calibration wand. It is designed to be an integrated part of the wand and provide a uniform user experience. Figure 5 shows the physical form of WAVI. The device features a narrow cylinder at the bottom where the user can hold it. Space for vibration motors is provided in that part. WAVI then widens to house four lights that provide visual feedback. The lights are mirrored on two sides of the device to create an impression of circles when the wand is rotated. We chose green as the colour for the lights as it is bright and traditionally associated with positive results in computer systems. The handle form-factor allows for smooth rotation and the larger components are moved up to allow for a comfortable grip even for users with relatively large palms.

We decided to use a low number of lights so that the feedback is easy to perceive at a high rotation speed. Four lights light up when the volume in the immediate vicinity (approximately) of the wand is sufficiently calibrated. No lights are visible when the volume is uncalibrated and intermediate levels are also visualised. Haptic feedback is provided in the form of vibration stimulating the palm of the user's hand. Vibration motors operate on full power when the volume is not calibrated. No vibration is present for a fully calibrated volume. Intermediate levels are mapped to motor power levels in between. Figure 3 shows the relationship between calibration levels and provided feedback. WAVI does not change the calibration procedure in any way. The manufacturer's instructions fully apply and the procedure is initiated, conducted and assessed identically to calibrating without the use of WAVI. A final design problem to be solved was tracking the calibration procedure in order to determine what feedback should be provided with WAVI at any given time.

Modelling the calibration progress

We decided to model the measurement space that needs to be calibrated as a rectangular cuboid. This model is derived from the observation that measurement volumes in motion capture are usually reported in terms of length × width × height [12]. A commonly stated guideline by optical motion capture system vendors is to cover the measurement volume evenly during calibration with the wand [29, 37].

To estimate coverage, we create a histogram of point counts over the cuboid, subdividing the cuboid into 3D bins of equal size. The binning permits a computationally tractable approximation of the distribution of tracked wand markers across the volume in real-time. We set a threshold on the bin's point count above which the bin is considered calibrated. Because WAVI feedback is designed to be mostly based on coverage, we set this threshold to 1.

To enable smoother transition from an uncalibrated to a calibrated state, and because the individual bins can be small, we opted to consider neighboring bins in the calculation of the calibration progress. This also enables WAVI to provide feedback about not only the bins in which the wand markers are placed, but also the vicinity of the wand. For this purpose, we introduce *calibration levels*. We define the uncalibrated level to be 0 and the maximum level to be an integer parameter C_{maxlevel}. We then calculate the calibration level of a particular bin v as

$$\text{level}(v) = \text{round}\left(C_{\text{maxlevel}} \cdot \frac{\mathbb{1}_{h(v)>0} + \sum_{w \in \mathcal{N}_v} \mathbb{1}_{h(w)>0}}{|v \cup \mathcal{N}_v|}\right) \quad (1)$$

Here, round(·) is the rounding function (round to nearest integer), $h(\cdot)$ denotes the count of the respective bin, $\mathbb{1}_{h(\cdot)>0}$ is the indicator function that is 1 if $h(\cdot) > 0$, and 0 otherwise, $\mathcal{N}_v$ is the set of the (twenty six) neighbours of bin v, and $|\cdot|$ is the set cardinality (twenty seven). Since the wand has two markers and thus its spatial position is described by two bins, we calculate the current calibration level in WAVI's vicinity as the average of the calibration levels of the respective bins. This information can be translated to haptic and visual feedback.

IMPLEMENTATION

WAVI was implemented as a proof-of-concept prototype consisting of 3 components (Figure 4): (i) the WAVI device, (ii) a Qualisys Oqus motion capture system, and (iii) the WAVI Controller application. The individual components communicate constantly during the calibration process. Additionally, we built analysis tools for the purpose of evaluation.

The WAVI device

The user interaction in WAVI is handled through a custom-built device mounted on the calibration wand. The device is an ARM-based Arudino prototype built with the RFDuino development kit. WAVI is fully reprogrammable, so refinements to feedback mapping can be easily introduced. The integrated Bluetooth Low Energy (BLE) chip is responsible for communication with other parts of the system.

Figure 5 shows the physical layout of the WAVI device. Two standard mobile phone 2.5 $mm \times 10\, mm$ disc-shaped vibration

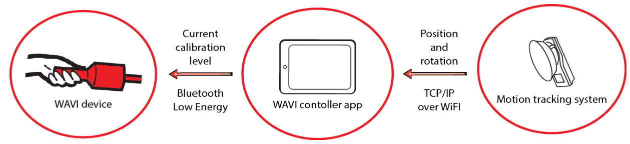

Figure 4: The three main components of the WAVI systems and the communication between them. Positional data is obtained from the motion capture system and processed by the controller application which provides the WAVI device with the feedback mode required at a given time.

motors with a maximum speed of 12000±2500 RPM provide haptic feedback on WAVI's handle. The motors are glued directly onto the wand to assure effective propagation of the vibration. Two sets of four high powered green LEDs provide visual feedback. The LEDs are mounted on an inner chassis to ensure stability. Two AAA-sized batteries provide power for more than 80 calibrations over three days. A custom 3D-printed body protects the electronics and allows for easy disassembly so that the batteries can be replaced. We painted WAVI black to harmonise with the colour of the wand.

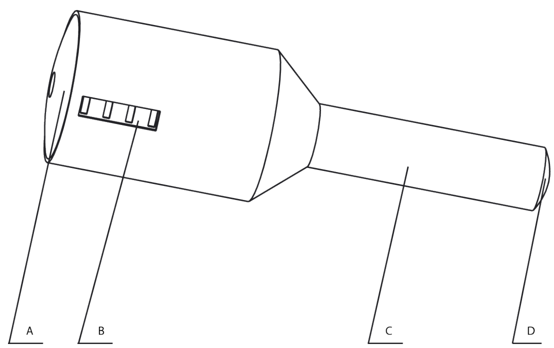

(a) Technical drawing of the 3D printed enclosure. A – top cap, B – inner chassis with LED mounts, C – handle part with separation between the wand and the enclosure to fit vibration motors, D – bottom cap.

(b) WAVI in action. Three of four LEDs are activated, showing a calibration level of approximately 75%.

Figure 5: The physical design of WAVI. Concept (a) and implementation (b).

Qualisys Oqus motion capture system

The motion capture system consists of infrared sensing cameras and the Qualisys Track Manager (QTM) software. The cameras are connected to a local network via Ethernet. A computer running QTM is connected to the same network. QTM reconstructs the 3D locations from the marker data captured by the cameras. The 3D marker positions can be queried in real-time by connecting to QTM via TCP or UDP. To be able to connect to QTM from other applications, we created a Wifi network and linked the computer running QTM to it.

WAVI Controller application

The WAVI Controller application is an Android-based software that tracks the calibration progress and sends control messages to the WAVI device. It connects to QTM via TCP and to the WAVI device via BLE. The application enables the operator to engage and disengage feedback and manually set WAVI to any state. These features facilitate conducting experiments with the device. We deployed the application on a Nexus9 tablet and connected the tablet via WiFi to the same network the QTM computer was connected to. The reason for deploying the controller application on a tablet was to facilitate the experimental setup for the evaluation. The application could, however, be deployed on the same computer as QTM, thus not requiring additional external hardware.

Real-time data processing

From QTM the WAVI Controller application requests the 3D marker data every 100 ms. The application then identifies the two points corresponding to the wand markers based on the known distance between the two markers. The wand marker points are added to their respective bins in the cuboid. The identification of the wand markers may fail, for example because only one or no marker was visible by the camera system. In that case, the processing is aborted for the current sample. After incrementing the bin count, the calibration level is calculated according to Equation 1. We determined through informal user studies that 5 levels $(0, 1, \ldots 4)$ are appropriate, and thus set the maximum level, C_{maxlevel}, to 4. The calculated level is sent as a control signal to the WAVI device. There, the calibration level is mapped to the corresponding number of LEDs lighting up and to one of 5 discrete linearly-distributed pulse-width modulation signals for the vibration motors (maximum vibration at level 0, no vibration at level 4).

For the purpose of post-hoc data analysis, we built a calibration density visualising tool in Mathematica. This tool can produce a heat map of the calibrated volume if provided with a recording of a calibration. The tool enabled us to later investigate the calibration patterns produced by users.

EVALUATION

We conducted a controlled within-subject experiment where we compared calibration using WAVI with standard calibration without feedback. Each subject calibrated the volume under two conditions: with WAVI activated (**W**) and without any feedback (**NW**). We investigated the following hypotheses: **[H1]** Using WAVI will result in better calibration quality than standard calibration; **[H2]** Calibration with WAVI will be perceived as requiring less effort than standard calibration; **[H3]** Users will prefer using WAVI compared to calibration without WAVI.

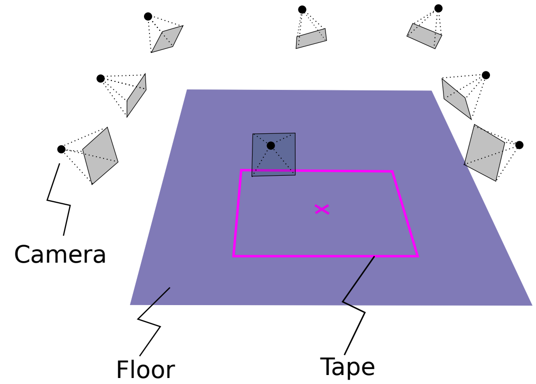

Figure 6: Experimental set-up: An eight-camera configuration suspended on a ceiling-mounted scaffold was used. The cameras were placed at heights 2.9 *m* to 3.5 *m* above the floor and were vertically inclined. The L frame was fixed to the back right corner.

Participants

We recruited 20 participants (14 male, 6 female, aged between 18 and 50 years, $\mu = 29.9$, $\sigma = 9.28$) through word of mouth. Each participant was remunerated with a gift of their choice at a maximum value of USD 20. Participants were unaware of the purpose of the experiment prior to the experimental session. Each participant attended one experimental session lasting for a maximum 30 *min*. duration. All participants reported that they did not have any medical conditions restricting their movement. We also recorded the height of each participant to later check for possible bias introduced by differences in arm length. All participants had not calibrated a marker-based motion tracking system prior to the study. Users were randomly assigned a starting condition and order balancing was applied so that 10 users began with the NW condition followed by W; 10 followed the opposite order.

Apparatus

The experiment was conducted in a dedicated windowless motion tracking room with controlled lighting. An illustration of the experimental setup is shown in Figure 6. The space to be calibrated was empty and had dimensions of $3.5\,m \times 3.5\,m \times 2.75\,m$ (length × width × height), with ceiling height exceeding 3.5 *m*. We outlined the edges of a square of 3.5 *m* side length on the floor using pink adhesive tape. Additional space of more than 2 *m* was available on each side of the square. We marked the center of the square on the floor with a taped "X" sign.

To record calibration trials and obtain the wand marker positions in real-time, an eight-camera Qualisys Oqus 500+ system was used. Each camera was equipped with a 13 *mm* lens and had 2048 × 2048 image resolution. The cameras were suspended in a rectangular configuration on a ceiling-mounted scaffold, placed between 2.9 *m* and 3.5 *m* above the floor and vertically inclined so that the floor-area outlined by the pink tape was centered in the respective camera's view.

The participants used a wand with WAVI attached and two markers 500.9 *mm* apart; the L-frame was placed on the floor on one of the corners of the square outlined by the tape. To be able to analyze trajectory data from recorded calibration trials, the system was calibrated by an experienced user. This enabled us to easily record participant calibration and process the measurements off-line for further analysis. We used QTM to record the trials, the reference calibration and the three-dimensional trajectory reconstruction. Recordings were made at 60 frames per second.

To track the calibration progress and calculate the calibration level (Equation 1), we chose a bin size of $350\,mm$ in each dimension for the cuboid subdivision. This resulted in a histogram of $10 \times 10 \times 7$ bins (length $\times$ width $\times$ height). The selected bin size was smaller than the distance between the markers on the wand, ensuring the two markers would not fall into the same bin at the same time. By not making the bins smaller, we made it easier to achieve a higher overall coverage of the volume.

Measures

A subset of dependent variables to be measured was derived from the output produced by QTM after each calibration trial:

Points obtained: the number of wand marker points captured, per camera and calibration. The more points, the better the calibration. QTM reports the point counts per camera, and we compute the sum over all cameras.

Standard deviation of wand length: a measure of how much the wand length modelled in 3D varies compared to the known length, per calibration. A lower standard deviation indicates better calibration. Measured in millimetres [*mm*]. We use the value as reported by QTM.

Residuals: the sum of average triangulation errors within the volume where the wand was placed, per camera and calibration. A lower residual corresponds to better calibration. Measured in millimetres [*mm*]. QTM reports the average residuals per camera, and we calculate the sum over all cameras.

Volume coverage: percentage of the target volume covered by the participant, per calibration. This is an estimate of how evenly the wand points were distributed over the target volume. A higher coverage corresponds to better calibration. QTM does not report this value.

We re-processed the recorded trials in QTM to obtain the individual measures. As QTM does not report volume coverage, we exported the 3D wand marker data of each trial to a Matlab MAT file. We then used a custom Matlab script to recreate the same cuboid division used in the WAVI system, assign the 3D points to the respective bins, and finally calculate the volume coverage as the percentage of bins with point count > 0.

Additionally, we measured **perceived workload** by means of the "Raw" (without pairwise comparisons) NASA Task Load Index (RTLX, referred to as NASA TLX in the rest of this paper) questionnaire. We also asked users for **preference** between the W and NW conditions. A 20-point Likert scale was used to subjectively asses the **usefulness of visual and haptic feedback** (from 0 — highly unnecessary, to 20 — extremely useful). We also conducted semi-structured interviews with the participants to learn more about their experience with the prototype. After the experiment we analyzed **volume fill patterns** using our visualisation tool.

Experimental Procedure

The study began with each participant filling in a demographic questionnaire which also asked about their height and whether they had any conditions restricting movement. Next, the participant was presented with video instructions for the study. We opted for video instructions as this eliminated any variability in the instructions given. As a well-defined description of the task cannot be found in the literature, we analyzed the system manual [29] for any clues as to how to perform calibration. All that information was included in the video. The video gave examples of the use of motion tracking and its operating principle. It then described the calibration volume and the reference objects used. Instructions for the calibration task were given using the wording from the manufacturers manual. The participants were also instructed to cover the entire target volume. An additional video explained how WAVI represented calibration and was shown before the participant calibrated the volume in the W condition.

Next, the participant was asked to perform five calibration trials of 45 seconds in the first condition assigned. The chosen trial duration is a 50% increase from the vendor-recommended minimum for a "normal measurement volume" (20 — 30 seconds). We consider the extra time margin to be sufficient for the participants to achieve a correct calibration. At the beginning of each trial, the participant was instructed to step on the x-marked center of the outlined square, and wait for the experimenter to count down from five before starting the calibration trial. The NASA TLX questionnaire was administered after 5 trials. Five more trials in the other condition were then performed and the participant filled a second TLX form and a preference questionnaire. Participants were later debriefed with a semi-structured interview about their experiences and received remuneration. The procedure generated a total of 200 calibration and 40 TLX measures.

RESULTS

Calibration quality

First, we investigated whether condition order in the experiment had an effect on the measures. We performed a two-way ANOVA to test for order effects with the starting condition as a between-subject factor. No effect was observed. Next, we investigated the performance measures and performed one-way ANOVAs to investigate the effects. The performance results are shown in Figure 7.

The participants generated more points in the cameras when using WAVI. The mean number of points collected using WAVI per participant was μ = 3092, which is 1% more than the mean of μ = 3058 in the NW condition. The difference is statistically significant ($F_{(1,198)} = 19.96$, $p < 0.001$). Similarly, camera residuals were lower. The mean residual per camera, participant and calibration observed when using WAVI was μ = 1.160 mm, which is 2% less than the mean of μ = 1.185 mm obtained when the users were not using WAVI.

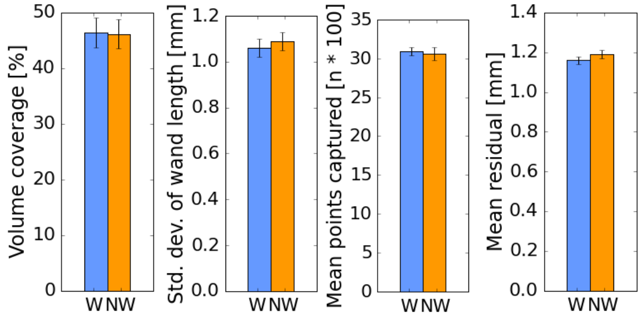

(a) (b) (c) (d)

Figure 7: Mean results for the calibration performance measures in the experiment. The mean difference is statistically significant for number of points collected and camera residuals. Note that in case of standard deviation of wand length and camera residuals lower results indicate better performance. Error bars show standard deviations.

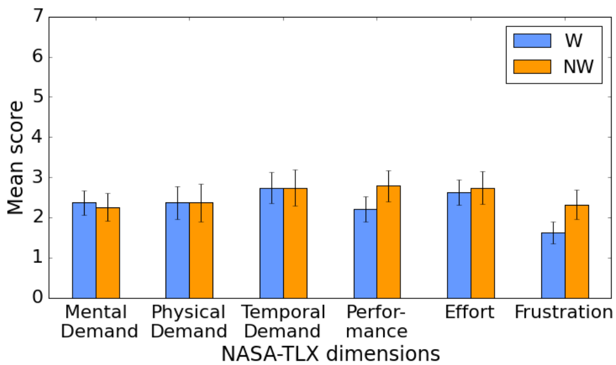

Figure 8: Mean NASA TLX questionnaire results for WAVI evaluation. The result is statistically significant for frustration. Note that we used NASA TLX a scale from 0 to 7 for the purposes of presentation. Lower scores indicate lower workload. Error bars show standard errors.

The difference is statistically significant with $F_{(1,198)} = 16.61$ and $p < 0.001$. While mean differences for the standard deviation of wand length and the volume covered were observed, no significant effect was present for either measure.

Perceived workload

NASA TLX results show that WAVI did not affect perceived workload in most TLX categories. We performed a Wilcoxon signed-rank test for those categories where a mean difference was present. A significant effect on frustration was observed with $p < 0.05$ (corrected for ties). Figure 8 shows the results of the assessment.

Preference, feedback modalities

Overall, 17 out of 20 users preferred using WAVI to not using the device (median $\tilde{x} = 3$ with 0 — maximum preference for W and 20 — maximum preference for NW). Vibration feedback was perceived as very useful ($\tilde{x} = 20$ with $x_{min} = 7$) while visual feedback was subjectively assessed as moderately beneficial to the task ($\tilde{x} = 9$).

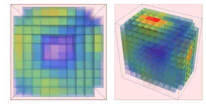

(a) Uniform distribution.

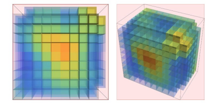

(b) Mostly in the central zone.

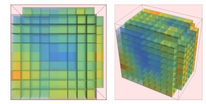

(c) Mostly near the edges.

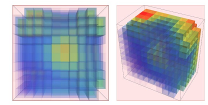

(d) Localised.

Figure 9: Point concentration histograms for the four user groups identified in our study based on the volume fill patterns generated during calibration.

Volume fill patterns

Using the custom-built Mathematica tool, volume fill patterns were visualised for each participant and condition thus creating 40 visualisations of how users filled the calibration volume. Figure 9 shows an example of how a user's volume fill pattern was visualised. We then endeavored to identify patterns and user groups within the data. Three researchers participated in a session of iterative open coding where patterns in the visualisations were identified. We then invited three different researchers to verify the patterns. We also noted that 19 out of 20 entered each bin at least once. The analysis resulted in discerning four distinct user groups:

The first group are users who achieved a *uniform distribution of calibration points* (Figure 9a). In 6 cases, we identified that the distribution pattern was sufficient for calibrating the volume and the users did not omit any areas of the volume. The patterns generated in 12 cases the majority of points fell within the *central zone of the calibration volume* (Figure 9b). These users spent the majority of their calibration time around the X-marked spot on the floor of the calibration volume and rarely entered zones located closer to the edges of the calibration volume. In contrast, five cases exhibited a *higher concentration of points near the edges of the volume* (Figure 9c). These patterns indicate a user focused on the boundaries of the volume, which resulted in lower quality calibration in the central zones. Finally, 17 patterns were *localised around a point or edge* (Figure 9d). Users would pick one corner or edge of the calibration volume and spend most of the calibration time around that zone.

DISCUSSION

The results of the study support hypothesis **H1** and show that providing feedback with WAVI did result in significant improvements in calibration quality (most importantly the number of points collected and camera residuals). This implies that the design of WAVI was successful and future systems should incorporate in-calibration feedback. Clearly, the status quo (i.e. feedback provided only at the end of calibration procedure) generates less-than-optimal results. This is why we believe WAVI has effectively demonstrated that in-calibration feedback is an effective way of improving calibration quality, previously not identified in motion capture research.

Furthermore, hypothesis **H2** was partially confirmed as the study indicates that using WAVI reduced users' frustration. We believe this result stems from the fact that novel users have little knowledge of how a well-calibrated space impacts measurement and almost no expectations as to the results of the calibration. Unlike expert users, they are not accustomed to calibrating and checking for calibration quality afterwards. A reasonable explanation would be that WAVI offers continuous feedback throughout the calibration and users are less frustrated as they have a better perception of completing the task. In other words, the changes in feedback enable users to perceive task completion, while the task in the NW condition is more ephemeral. Having said that, we were surprised not to observe a significant effect on perceived performance (in the NASA TLX) as WAVI also provided immediate information on the calibration level. The lack of significant difference in mental demand shows that calibrating WAVI is perceived as equally complex as just using the calibration wand. The preference questionnaire support hypothesis **H3**; users strongly preferred a device that offered feedback to one with no feedback even if their subjective perception of the task was not significantly altered by the device.

We also observed that haptic feedback proved to be effective and it was well received by the users. While we did find a solid rationale for using haptic feedback in 3D input in related work, we were unsure whether these findings would be applicable to a calibration task. The intricate random movements of the wand dance are quite different from, for instance, 3D painting. Our results show that the use of vibratory output may be appropriate in a wide range of 3D input tasks.

Our results also show that many users skipped certain parts of the calibration volume. In a measurement setting, this would result in decreased tracking accuracy in the zones with less calibration. As we observed a significant increase in volume coverage when using WAVI, we can theorise that the perception of the calibration level in the vicinity of the wand allowed users to direct their attention to other parts of the calibrated volume. Consequently, more parts of the target cuboid were entered by the wand, resulting in improved volume coverage. Our analysis of volume fill patterns indicates that some users have a tendency to concentrate on particular parts of the volume. A more uniform distribution of points would result in better calibration quality. This indicates an emergent user need for providing enough information to users so that they are aware that parts of the volume are still insufficiently calibrated.

Future directions for calibration feedback

As we recognise that our work is the first inquiry into in-calibration feedback for motion capture, we believe that our inquiry yields new results in terms of issues that need to be explored further. Our work shows that providing feedback is beneficial to the user, and the calibration quality, the exact factors, and design constraints contributing to that effect are still unexplored. In this section, we identify key directions for further inquiry in developing systems with design goals similar to WAVI.

Firstly, we see an emergent need for designing *the means to provide directional feedback*. Our study shows that users often skip parts of the target volumes and directing them in 3D space so that the point count is more uniform would alleviate that issue. Finding effective measures to achieve that may be difficult as the position of the user is unknown to the system before tracking is configured. We believe that algorithms that use the approximate position of the user or using multiple sensors may be a solution to this issue. A remaining question is what modalities to use to provide directional guidance and how to achieve that without considerable investments in infrastructure. The user groups that we identified in the study indicate that many users would benefit from directing them during calibration. The high number of localised patterns shows that an increased perception of directionality could improve calibration quality.

Secondly, the way users performed given a well-defined volume in our study suggests that future systems should address the challenge of providing the means for *clearer volume definitions*. In a field measurement setting, the user usually knows the target calibration volume. It is vital for the motion capture system offering the desired accuracy that the target volume is entered with the wand. Current systems simply assume that the boundaries of the volume defined using the wand dance procedure are the boundaries of the target calibration volume. Our work shows that there is an emergent need to compare the volume the user desires to calibrate with the one actually calibrated. We believe that volume definition procedures can be designed and integrated in motion capture systems. In-calibration feedback based on the volume definition could indicate to the user the overall progress. For instance, the volume coverage measure as presented in this paper could be one possible progress indicator. We only considered feedback based on location-based progress measures in this work. On the other hand, we cannot dismiss that indicating temporal progress of the calibration procedure to users might have a positive impact on calibration quality.

Thirdly, as our work explored only local feedback (i.e. concerning the volume in the immediate vicinity of the wand as opposed to the whole target volume), we believe that future designs should investigate *the role of global feedback and the temporal aspect of calibration*. As we observed that many users found it difficult to provide uniform input throughout the target volume within the set calibration time, it is meaningful to investigate ways for users to pace their calibration well. Differences in the perception of time are also a possible explanation for the variety of volume fill patterns that we observed. Given that many patterns were localised, it is possible that some users decided to focus on one zone with the intention to attend to other zones later, but they ran out of calibration time. While our study showed local feedback may lead to improvements in calibration quality, future work should explore whether additional global feedback can lead to further quality improvements.

Lastly, as WAVI illustrates that vibratory feedback can be an effective means of providing in-calibration output, the logical next step is to *explore different designs of haptic feedback* for motion capture calibration. While there is extensive research on designing haptic output in HCI, our study only employed basic vibratory stimulus as it was the first inquiry of its kind. We see a need for reapplying the lessons learned from applications such as mobile settings [3] or personal navigation [27] in dynamic 3D input tasks. This will enable future calibration support systems to communicate more accurate information to users and achieve better calibration performance.

Limitations

As our work was the first inquiry into in-calibration feedback, it is constrained by certain limitations. Our study was limited to a single motion tracking system from a single manufacturer. While most marker-based motion tracking systems on the market follow a similar principle of operation, differences between models should be investigated in the future. This includes a comparison of existing feedback mechanisms (e.g. visual feedback on stationary screen, and on-camera visual feedback) with each other and with WAVI. It should be pointed out that our study was conducted with novice users and the overall calibration quality was rather low in absolute terms (on average, the users did not manage to cover even half of the target volume). It remains to be investigated what the effect of feedback would be on expert users. Such studies, however, will be quite hard to execute as the expert user group is rather limited in size. Another aspect that we did not explore is long-term usage. It is possible that users preferred using WAVI because of the aesthetic appeal of the device and not the improvements in the perception of the tasks. We see a need for studying how users appropriate calibration technology over longer periods of time.

As we endeavoured to identify whether feedback would produce any positive results in our inquiry, we looked for the effects of using WAVI on calibration quality. However, it may be possible that in-calibration feedback has an effect on how fast users learn to calibrate volumes sufficiently. Answering this question would require an entirely different experimental design, which we see as an emerging question for future research. Finally, we recognise that the lack of prior work in the area caused the need to make some design decisions based on our personal experiences and design intuition. While this is an essential step in an interaction design process, we see that the results of the study may be affected by our design decisions. We hope that future research will produce more in-calibration device prototypes and enable us to get a better understanding of their design through comparison.

The WAVI system uses positional data about the wand during calibration, which requires a calibrated motion capture system in the first place. We used a pre-calibrated system for this study. However, in real-world settings the motion capture system is initially uncalibrated. This bootstrap problem can be mitigated by, for instance, performing a quick initial calibration in a significantly smaller volume, implementing an on-the-fly initial calibration based on the first few frames of the calibration measurement [1, 17], or using a secondary, relative-position tracking system during calibration (e.g. based on inertial measurement unit (IMU) technology). Since we were primarily interested in the effects of the feedback provided by the WAVI system, and to reduce additional variability introduced by a particular choice of bootstrapping strategy, we assumed an ideal setting where we could rely on accurate positional data to test our hypotheses.

CONCLUSIONS

In this paper we presented the design, implementation and evaluation of WAVI — a proof-of-concept prototype of an in-calibration feedback device for motion capture calibration. WAVI provides visual and haptic feedback during calibration to help users achieve better calibration quality. In a within-subjects experiment, we showed that WAVI significantly improved calibration quality and reduced user frustration. Users preferred using WAVI compared to lack of feedback and the haptic feedback was perceived as more useful. Based on the results of the study, we show that in-calibration feedback has the potential to improve the quality of working with motion capture systems. We also analyzed the volume fill patterns generated by the participants to identify user groups among them. We conclude with a set of emerging challenges for designing calibration procedure for motion tracking. As our work is the first to investigate motion tracking calibration in the field of HCI, we hope that this paper will inspire further inquiries that will result in an enhanced understanding of the complex input tasks involved in motion capture calibration thus contributing to designing more efficient tools for motion capture systems.

ACKNOWLEDGEMENTS

This work was supported by the EU FP7 People Programme (REA Grant Agreements 289404 and 290227), and by the Adlerbertska Research Foundation.

The authors would like to thank Igor Stanković for implementing an initial version of the Mathematica tool, Ayça Ünlüer for designing some of the figures used in this paper and Philippa Beckman for proofreading.

REFERENCES

1. Motilal Agrawal and Larry S. Davis. 2003. Camera calibration using spheres: A semi-definite programming approach. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*. IEEE, 782–789.

2. N. Alberto Borghese and Pietro Cerveri. 2000. Calibrating a video camera pair with a rigid bar. *Pattern Recognition* 33, 1 (2000), 81–95.

3. Jessalyn Alvina, Shengdong Zhao, Simon T. Perrault, Maryam Azh, Thijs Roumen, and Morten Fjeld. 2015. OmniVib: Towards Cross-body Spatiotemporal Vibrotactile Notifications for Mobile Phones. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 2487–2496. DOI: http://dx.doi.org/10.1145/2702123.2702341

4. Myroslav Bachynskyi, Gregorio Palmas, Antti Oulasvirta, Jürgen Steimle, and Tino Weinkauf. 2015.

Performance and Ergonomics of Touch Surfaces: A Comparative Study using Biomechanical Simulation. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1817–1826.

5. Seok-Hyung Bae, Ravin Balakrishnan, and Karan Singh. 2008. ILoveSketch: As-natural-as-possible Sketching System for Creating 3D Curve Models. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology (UIST '08)*. ACM, New York, NY, USA, 151–160. DOI: http://dx.doi.org/10.1145/1449715.1449740
6. Elena Ceseracciu, Zimi Sawacha, and Claudio Cobelli. 2014. Comparison of markerless and marker-based motion capture technologies through simultaneous data collection during gait: proof of concept. *PloS one* 9, 3 (2014), e87640.
7. Veronika Fedotova and Viktors Piipivs. 2012. 3D Camera Calibration for Motion Tracking in Recurrent Athletic Environment. In *5th Baltic Sport Science Conference Current Issues and New Ideas in Sport Science*. 64. https://ortus.rtu.lv/science/en/publications/13111-3D+Camera+Calibration+for+Motion+Tracking+in+Recurrent+Athletic+Environment
8. David R. Flatla, Carl Gutwin, Lennart E. Nacke, Scott Bateman, and Regan L. Mandryk. 2011. Calibration Games: Making Calibration Tasks Enjoyable by Adding Motivating Game Elements. In *Proceedings of UIST '11*. ACM, 403–412. DOI: http://dx.doi.org/10.1145/2047196.2047248
9. Jonas Fredriksson, Sven Berg Ryen, and Morten Fjeld. 2008. Real-time 3D hand-computer interaction. In *Proceedings of the 5th Nordic conference on Human-computer interaction building bridges - NordiCHI '08*. ACM Press, New York, New York, USA, 133. http://dl.acm.org/citation.cfm?id=1463160.1463175
10. A.D. Gregory, S.A. Ehmann, and M.C. Lin. 2000. inTouch: interactive multiresolution modeling and 3D painting with a haptic interface. In *Proceedings IEEE Virtual Reality 2000 (Cat. No.00CB37048)*. IEEE Comput. Soc, 45–52. http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=840362
11. Tovi Grossman and Ravin Balakrishnan. 2004. Pointing at Trivariate Targets in 3D Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 447–454. DOI: http://dx.doi.org/10.1145/985692.985749
12. Gutemberg Guerra-Filho. 2005. Optical Motion Capture: Theory and Implementation. *RITA* 12, 2 (2005), 61–90.
13. Huirong Han, Juli Yamashita, and Issei Fujishiro. 2002. 3D Haptic Shape Perception Using a 2D Device. In *ACM SIGGRAPH 2002 Conference Abstracts and Applications (SIGGRAPH '02)*. ACM, New York, NY, USA, 135–135. DOI: http://dx.doi.org/10.1145/1242073.1242148
14. Thomas Hermann, Andy Hunt, and John G. Neuhoff. 2011. *The sonification handbook*. Logos Verlag Berlin.
15. Otmar Hilliges, David Kim, Shahram Izadi, Malte Weiss, and Andrew Wilson. 2012. HoloDesk: Direct 3D Interactions with a Situated See-through Display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 2421–2430. DOI: http://dx.doi.org/10.1145/2207676.2208405
16. Martin Hoßbach, Johannes Gregori, Stefan Wesarg, and Matthias Günther. 2013. Design and analysis of a calibration-method for stereo-optical motion tracking in MRI using a virtual calibration phantom. In *SPIE Medical Imaging*, Robert M. Nishikawa and Bruce R. Whiting (Eds.). International Society for Optics and Photonics, 86682E. http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=1661963
17. Zuzana Kukelova, Martin Bujnak, and Tomas Pajdla. 2013. Real-time solution to the absolute pose problem with unknown radial distortion and focal length. In *Computer Vision (ICCV), 2013 IEEE International Conference on*. IEEE, 2816–2823.
18. Jonas Löwgren. 2013. Annotated Portfolios and Other Forms of Intermediate-level Knowledge. *interactions* 20, 1 (Jan. 2013), 30–34. DOI: http://dx.doi.org/10.1145/2405716.2405725
19. Pengfei Lu and Matt Huenerfauth. 2009. Accessible Motion-capture Glove Calibration Protocol for Recording Sign Language Data from Deaf Subjects. In *Proceedings of the 11th International ACM SIGACCESS Conference on Computers and Accessibility (Assets '09)*. ACM, New York, NY, USA, 83–90. DOI: http://dx.doi.org/10.1145/1639642.1639658
20. S. McMains. 2004. Evaluation of drawing on 3D surfaces with haptics. *IEEE Computer Graphics and Applications* 24, 6 (Nov. 2004), 40–50. http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1355891
21. Stephen W. Mereu and Rick Kazman. 1996. Audio enhanced 3D interfaces for visually impaired users. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 72–78.
22. Thomas B. Moeslund and Erik Granum. 2001. A survey of computer vision-based human motion capture. *Computer vision and image understanding* 81, 3 (2001), 231–268.
23. Thomas B. Moeslund, Adrian Hilton, and Volker Krüger. 2006. A survey of advances in vision-based human motion capture and analysis. *Computer vision and image understanding* 104, 2 (2006), 90–126.

24. NaturalPoint, Inc. 2015. Calibration - NaturalPoint Product Documentation. (2015). http://wiki.optitrack.com/index.php?title=Calibration

25. Matti Niinimäki and Koray Tahiroglu. 2012. AHNE: a novel interface for spatial interaction. In *Proceedings of the 2012 ACM annual conference extended abstracts on Human Factors in Computing Systems Extended Abstracts - CHI EA '12*. ACM Press, New York, New York, USA, 1031. DOI: http://dx.doi.org/10.1145/2212776.2212378

26. Ken Pfeuffer, Melodie Vidal, Jayson Turner, Andreas Bulling, and Hans Gellersen. 2013. Pursuit calibration: making gaze calibration less tedious and more flexible. In *Proceedings of the 26th annual ACM symposium on User interface software and technology - UIST '13*. ACM Press, New York, New York, USA, 261–270. DOI: http://dx.doi.org/10.1145/2501988.2501998

27. Manoj Prasad, Paul Taele, Daniel Goldberg, and Tracy A. Hammond. 2014. HaptiMoto: Turn-by-turn Haptic Route Guidance Interface for Motorcyclists. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3597–3606. DOI: http://dx.doi.org/10.1145/2556288.2557404

28. Tomislav Pribanić, Peter Sturm, and Mario Cifrek. 2007. Calibration of 3D kinematic systems using orthogonality constraints. *Machine Vision and Applications* 18, 6 (2007), 367–381.

29. Qualisys AB. 2015. Qualisys Track Manager (QTM) user manual. (2015).

30. Victoria Rudakova and Pascal Monasse. 2014. Camera matrix calibration using circular control points and separate correction of the geometric distortion field. In *Computer and Robot Vision (CRV), 2014 Canadian Conference on*. IEEE, 195–202.

31. Steven Schkolne, Michael Pruett, and Peter Schröder. 2001. Surface Drawing: Creating Organic 3D Shapes with the Hand and Tangible Tools. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*. ACM, New York, NY, USA, 261–268. DOI: http://dx.doi.org/10.1145/365024.365114

32. Ki-Young Shin and Joung Hwan Mun. 2012. A multi-camera calibration method using a 3-axis frame and wand. *International Journal of Precision Engineering and Manufacturing* 13, 2 (2012), 283–289.

33. Zhongwei Tang, Rafael Grompone von Gioi, Pascal Monasse, and Jean-Michel Morel. 2012. High-precision camera distortion measurements with a "calibration harp". *JOSA A* 29, 10 (2012), 2134–2143.

34. Robert J. Teather and Wolfgang Stuerzlinger. 2013. Pointing at 3d target projections with one-eyed and stereo cursors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*. ACM Press, New York, New York, USA, 159. http://dl.acm.org/citation.cfm?id=2470654.2470677

35. Yu Ukai and Jun Rekimoto. 2013. Swimoid: A Swim Support System Using an Underwater Buddy Robot. In *Proceedings of the 4th Augmented Human International Conference (AH '13)*. ACM, New York, NY, USA, 170–177. DOI: http://dx.doi.org/10.1145/2459236.2459265

36. Janet van der Linden, Rose Johnson, Jon Bird, Yvonne Rogers, and Erwin Schoonderwaldt. 2011. Buzzing to Play: Lessons Learned from an in the Wild Study of Real-time Vibrotactile Feedback. In *Proceedings of CHI '11*. ACM, 533–542. DOI: http://dx.doi.org/10.1145/1978942.1979017

37. Vicon Motion Systems Limited. 2015. Vicon Tracker User Guide. (2015).

38. Arantxa Villanueva, Rafael Cabeza, and Sonia Porta. 2004. Eye Tracking System Model with Easy Calibration. In *Proceedings of the 2004 Symposium on Eye Tracking Research & Applications (ETRA '04)*. ACM, New York, NY, USA, 55–55. DOI: http://dx.doi.org/10.1145/968363.968372

39. Andrew Wilson and Steven Shafer. 2003. XWand: UI for Intelligent Spaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*. ACM, New York, NY, USA, 545–552. DOI: http://dx.doi.org/10.1145/642611.642706

40. Markus Windolf, Nils Götzen, and Michael Morlock. 2008. Systematic accuracy and precision analysis of video motion capturing systems–exemplified on the Vicon-460 system. *Journal of biomechanics* 41, 12 (Aug. 2008), 2776–80. http://www.sciencedirect.com/science/article/pii/S0021929008003229

41. Robert Xiao, Greg Lew, James Marsanico, Divya Hariharan, Scott Hudson, and Chris Harrison. 2014. Toffee: enabling ad hoc, around-device interaction with acoustic time-of-arrival correlation.. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services - MobileHCI '14*. ACM Press, New York, New York, USA, 67–76. DOI: http://dx.doi.org/10.1145/2628363.2628383