

# Precise Haptic Device Co-Location for Visuo-Haptic Augmented Reality

Ulrich Eck, *Student Member, IEEE*, Frieder Pankratz, Christian Sandor, *Member, IEEE*, Gudrun Klinker, *Member, IEEE*, and Hamid Laga

**Abstract**—Visuo-haptic augmented reality systems enable users to see and touch digital information that is embedded in the real world. PHANToM haptic devices are often employed to provide haptic feedback. Precise co-location of computer-generated graphics and the haptic stylus is necessary to provide a realistic user experience. Previous work has focused on calibration procedures that compensate the non-linear position error caused by inaccuracies in the joint angle sensors. In this article we present a more complete procedure that additionally compensates for errors in the gimbal sensors and improves position calibration. The proposed procedure further includes software-based temporal alignment of sensor data and a method for the estimation of a reference for position calibration, resulting in increased robustness against haptic device initialization and external tracker noise. We designed our procedure to require minimal user input to maximize usability. We conducted an extensive evaluation with two different PHANToMs, two different optical trackers, and a mechanical tracker. Compared to state-of-the-art calibration procedures, our approach significantly improves the co-location of the haptic stylus. This results in higher fidelity visual and haptic augmentations, which are crucial for fine-motor tasks in areas such as medical training simulators, assembly planning tools, or rapid prototyping applications.

**Index Terms**—Calibration, orientation calibration, temporal alignment

## 1 INTRODUCTION

RESEARCHERS have started to combine augmented reality (AR) and haptic interaction to enable users to see and touch digital information that is embedded in the real world. Such visuo-haptic augmented reality (VHAR) user interfaces with co-located visual augmentations and haptic interaction improve realism [1] and enable users to interact more precisely [2].

Haptic devices for providing kinesthetic feedback are usually based on one of the two concepts: stylus- and grip-based devices for tool interaction and string-based systems. Massie and Salisbury [3] developed the widely used stylus-based PHANToM device, which consists of two interlinked joints. The angles of these joints define the position of the gimbal, commonly called haptic interface point (HIP). The stylus can be rotated around the HIP while three sensors sense its orientation. The haptic stylus is often augmented with some context-dependent tool like a drill for dental surgery training [4], a brush for virtual painting [5], or tools for rapid prototyping [6]. Alternatively, the stylus can be hidden [7].

Precisely calibrating the components of a VHAR system (external trackers, cameras, haptic devices) and the spatial relations between them is essential to provide a realistic

user experience. Specifically, the integration of haptic devices is not trivial. The core challenge is to precisely determine the position and orientation of the haptic stylus, which are required in order to precisely co-locate visual augmentations and haptic feedback.

Various methods [8], [9], [10] have been proposed to integrate haptic devices into AR environments, but all of them focused on the calibration of the HIP position, which is required to co-locate linear force feedback. However, the stylus orientation is not considered. Fig. 1a shows a visual overlay with the current state-of-the-art position calibration and Fig. 1b shows a visual overlay with position and orientation calibration.

Improved position and orientation accuracy results in higher fidelity visual and haptic augmentations, which is crucial for fine-motor tasks in areas including medical training simulators, assembly planning tools, or rapid prototyping applications. A user friendly calibration procedure is essential for real-world applications of VHAR.

PHANToM devices need to be initialized every time they are powered on. Desktop and Premium models require users to manually place the device in a predefined reset pose without providing a mechanical fixture. Any deviation from the optimal reset pose will cause non-linear errors in the reported positions, since the assumptions in the forward kinematic model [11] do not match the physical state of the device. These errors are typically compensated in the position calibration. While using our platform however, we found that the current state-of-the-art calibration method proposed by Harders [10] is not able to achieve optimal results in all cases if no mechanical fixture is used during device initialization.

The position and orientation calibration needs reference measurements as inputs, which are used to calculate the

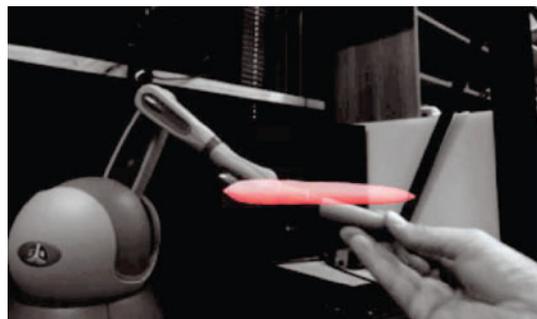
- U. Eck and H. Laga are with the Phenomics and Bioinformatics Research Centre, University of South Australia. E-mail: ueck@net-labs.de, hamid.laga@unisa.edu.au.
- F. Pankratz and G. Klinker are with the Fachgebiet Augmented Reality, Technische Universität München. E-mail: {pankratz, klinker}@in.tum.de.
- C. Sandor is with the Interactive Media Design Lab, Nara Institute of Science and Technology. E-mail: sandor@is.naist.jp.

Manuscript received 18 Jan. 2015; revised 6 July 2015; accepted 10 Sept. 2015. Date of publication 17 Sept. 2015; date of current version 26 Oct. 2015.

Recommended for acceptance by D. Schmalstieg.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TVCG.2015.2480087



(a) Virtual stylus overlaid with only joint angle calibration



(b) Overlay with our additional gimbal angle calibration.

Fig. 1. Visuo-haptic augmented reality system with a PHANToM haptic device. Previous calibration methods only compensate the errors in joint-angle sensors to improve position accuracy. Our method also calibrates the gimbal-angle sensors to reduce errors in the orientation of the stylus.

position and orientation errors. These measurements are typically provided by an external tracking system with sufficient accuracy. Precise temporal alignment of haptic poses and reference poses is essential for optimal calibration results. Previous work [10] used hardware synchronization to achieve temporal alignment, but not all workspace setups are compatible with this requirement.

In this article, we propose a complete calibration procedure for haptic devices in VHAR workspaces. The procedure calibrates both the position and orientation and thus it results in precise co-location of the haptic stylus and the associated haptic feedback. Our system design removes the requirement for hardware synchronization and minimizes the time and the amount of user input required for calibration. We conducted an extensive evaluation with two different PHANToMs, two different optical trackers, and a mechanical tracker. Our results show a significant reduction of the remaining position errors of up to 38 percent and the remaining orientation errors of up to 85 percent compared to the current state-of-the-art workspace calibration.

The contributions of this article and its associated conference paper [12] are fourfold. First, we introduce a novel algorithm to determine a reference for orientation calibration and a method for precisely calibrating the gimbal-angles. Second, we integrate a software-based approach to temporally align measurements between the external tracker and the haptic device, which enables continuous data collection for improved usability. Third, we extend Harders' method [10], which makes position calibration more robust against inaccurate haptic device initialization and external tracker noise. Finally, we designed our procedure to require minimal user input in order to maximize usability.

The rest of this article is structured as follows: Section 1.1 summarizes the related work. Section 1.2 gives an overview of the proposed framework and details the main contributions. Section 2 describes the procedure in detail and shows how it differs from Harders' method [10]. Section 3 presents an extensive evaluation. Finally, Section 4 concludes by discussing the generalizability of our results and presenting future work.

### 1.1 Related Work

Previous research has investigated the integration of haptic devices in AR systems. A key requirement is the precise co-location of the haptic stylus with the visual augmentations

to provide realistic force feedback and to enable precise visual overlays.

Vallino and Brown [8] pioneered the co-location of a PHANToM Omni haptic device in AR by determining an affine transformation between the haptic workspace and a world coordinate reference. They used four non-planar points, which are captured in both coordinate systems, to estimate the affine transformation. Their method does not take into account the fact that the poses reported by the haptic device are distorted due to errors in the joint angle sensors. Their method works only for static camera poses. Ikits et al. [9] presented a calibration method for haptic devices using a planar grid to correct the haptic stylus position. Their method significantly improves position accuracy in the region where the calibration grid was placed. However, in other areas of the workspace, this method fails to effectively compensate distortions. Wang et al. [13] derived a mathematical model of the co-location error for VHAR systems with half-mirror displays and defined error metrics. They propose a calibration method using their model and evaluate it using a precise mechanical tracker. Their approach results in improved co-location for half-mirror display setups, but ignores co-location errors caused by sensors in haptic devices.

The most important previous research is the work by Harders et al. [10]. They presented a calibration procedure for VHAR workspaces and an algorithm to compensate the non-linear distortion of position measurements of PHANToM haptic devices for the complete workspace. They used an external tracker and rigidly attached a tracking target to the haptic stylus to perform an open-loop calibration [14]. Their procedure is as follows: First, they determine the position of the HIP in relation to the tracking target using tooltip calibration as a reference. Second, they collect 30 position measurements reported from the haptic device and the corresponding tip positions from the tracker covering the complete workspace. Finally, they start an iterative optimization process, which consists of two altering steps. First, they determine the 6DOF transform between the external tracker and the haptic workspace, in order to express the reference positions in haptic device coordinates. Then, they compensate the joint angle sensor errors, which are causing the position errors. They model the sensor errors using a linear system and use Levenberg-Marquardt optimization to minimize the squared-distance between reference positions and positions calculated using the forward kinematic model of

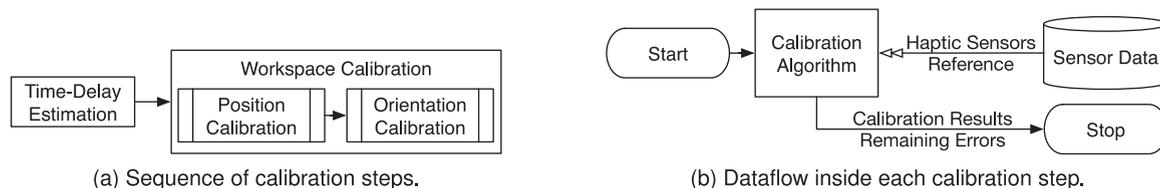


Fig. 2. Overview of the workspace calibration procedure. Initially, the delay between the external tracker and the haptic device is estimated. Then, position and orientation of the haptic stylus are calibrated. Each step consumes sensor data and returns calibration results and remaining errors.

the PHANToM from Çavuşoğlu et al. [11]. The iterative optimization is repeated until the change of the residual drops below a given threshold. Our calibration method additionally calibrates the orientation of the haptic stylus and further reduces position errors.

Knoerlein and Harders [15] compared tracker-based calibration with an alternative approach using a physical calibration object. They utilize known positional constraints of the haptic poses for optimization. During the calibration process, users slide the haptic stylus tip along the surface of three perpendicular planes to collect the samples for error compensation. By constraining the measured hip positions to these planes, they can effectively compensate position errors caused by the joint angle sensors. Their approach also compensates errors of the first two gimbal-angle sensors. Their method, however, only works for haptic devices with a physical tip, which is not the case for PHANToM Desktop and Premium 6DOF models.

All PHANToM haptic devices use digital encoders for sensing the joint angles. Gimbal angles are measured either by using digital encoders in Premium models, or by using analog potentiometers in Desktop and Omni models. The digital encoders used in PHANToM devices only report relative changes. Therefore, they need to be initialized every time the device is powered on. The initialization process varies between the models. PHANToM Omni devices provide an inkwell that holds the stylus in a predefined pose, which yields reproducible initialization. For all other PHANToM models, the user has to manually place the device into a reset pose. Any deviation from the optimal reset pose during initialization will cause non-linear position errors since the assumptions in the forward kinematic model [11] do not match with the physical state of the device.

Bianchi et al. [16] analyzed the effects of haptic device initialization quality in a simulation. They generated four data sets with varying initialization qualities. To enhance realism, they modeled the external tracker noise using a statistical noise model. The optimization process was carried out and results showed that the initial error could be reduced to mean values below 1.5 mm under all conditions. During the development of our platform however, we found that their iterative position calibration method does not yield optimal results for inaccurately initialized haptic devices. We extend their method with a novel initialization step to avoid any dependency on the haptic pose, which results in a significantly improved position and orientation calibration.

A key requirement for successful open-loop calibration is the precise temporal alignment of the corresponding measurements from the reference with the device that is calibrated. Hardware synchronization with a global trigger is

used in some environments [10] to ensure temporal alignment of sensor measurements. This is usually expensive and most devices do not directly support hardware synchronization without modifications. When unsynchronized sensors are used, a commonly used method to avoid inaccuracies is the point and hold method where corresponding samples are only recorded for static poses. This method is tedious and time-consuming. We utilize a software approach based on the work of Huber et al. [17] to temporally align measurements from unsynchronized sensors. This allows users to continuously move the haptic stylus during calibration, which results in improved coverage and faster calibration.

In our previous work [12], we presented a comprehensive workspace calibration method for accurate visuo-haptic augmented reality co-location, which extended the workspace calibration of Harders et al. [10]. We presented a novel algorithm for determining a reference orientation for the haptic stylus and a method for calibrating the gimbal angles. The integration of software-based sensor synchronization simplified calibration data acquisition and improved position and orientation calibration quality. The approach however has three limitations: First, the method did not yield optimal results for haptic devices without an inkwell. Second, the method required users to provide more input than necessary. Third, the method was not robust against tracker noise or dropouts. In this article we present an improved version of our calibration method [12], which further reduces the remaining position errors by up to 34 percent the remaining orientation errors by up to 86 percent and the required time for user input by more than 50 percent.

## 1.2 Overview and Contributions

In this article, we present a comprehensive workspace calibration method for accurate haptic device co-location in VHAR. An overview of our procedure is given in Fig. 2a. Our procedure is as follows:

- We estimate the time-delay between the external tracker and the haptic device using Huber's method [17] and compensate it. Subsequent steps in the procedure will receive temporally aligned measurements.
- We perform position calibration with an improved version of Harders' method [10] to calibrate the joint angles.
- We determine the orientation reference and use it to calibrate the three gimbal angles.

During each step, operators are moving the calibration target to provide the required input data. The calibration component receives sensor data streams from the haptic device and the tracking system. An error is calculated based

on a metric for each step. The calibration algorithm minimizes the remaining error after calibration by optimizing the parameters of the sensor error model. Each step yields a calibration result and the remaining errors after calibration. This process is shown in Fig. 2b.

The contributions of this article along with its earlier conference version [12] are as follows:

- Our procedure extends the current state-of-the-art workspace calibration with orientation calibration for accurate visual overlays and haptic co-location (see Figs. 1a and 1b). We introduce a novel algorithm for determining an orientation reference (Section 2.3.1) and a method to calibrate the gimbals angles (Section 2.3.2).
- We utilize a software-based approach for temporal alignment of unsynchronized sensor data in order to design a user friendly calibration procedure where, unlike previous methods [9], [10], users can continuously move the calibration target during data collection (Section 2.1).
- We minimize the required user input, streamline data collection and improve calibration-sample selection with the proposed data processing pipeline (Section 2.2).
- We improve position calibration by providing a better initial estimation of the transform between the external tracker and the haptic device (Section 2.2.1).

In Section 3, we present an extensive evaluation of various aspects of our calibration procedure in a low-fidelity setup (*LF*) and in a high-fidelity setup (*HF*). We also explore in a final verification the limits of our *HF* setup using a highly precise mechanical tracker. Our results show that calibrating gimbal sensors reduces the average orientation error by 63 percent for *LF* and 85 percent for *HF* in comparison to uncalibrated setups. We also show that the average position accuracy can be improved by up to 38 percent for *HF* in comparison to Harders' method [10], and that sub-millimeter accuracy can be achieved using our method.

## 2 CALIBRATION PROCEDURE

In this section, we provide details about our workspace calibration procedure and highlight the differences to existing methods. A schematic overview of our VHAR workspace setup is given in Fig. 3. The system is composed of an external IR tracker (*ET*), a camera (*C*), and a PHANToM haptic device (*HD*). The external tracker is used to track the camera viewpoint via its target ( $C_{target}$ ). The second target ( $HIP_{target}$ ) is rigidly attached to the stylus of the haptic device as reference for calibration. The spatial relation  $EH$  between the external tracker ( $ET_{origin}$ ) and the haptic workspace ( $HD_{origin}$ ) is static.

Initially, we calibrate the external tracker using the vendor-supplied tool, determine the camera intrinsics, and estimate the 6DOF transform between  $C_{target}$  and the camera coordinate system  $C$  using the hand-eye calibration method by Tsai and Lenz [18].

We initialize the PHANToM haptic device using the vendor-supplied tool and rigidly attach  $HIP_{target}$  to its stylus. The PHANToM Omni provides an inkwell for

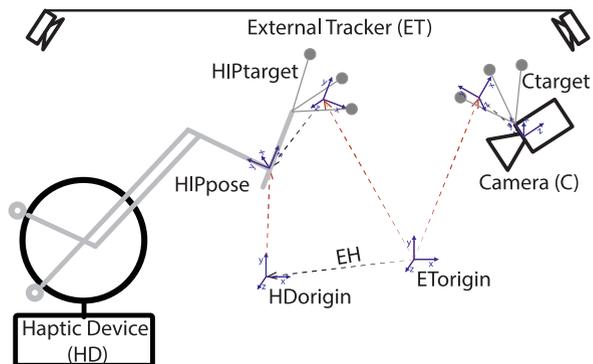


Fig. 3. Our setup for VHAR workspace calibration and its spatial relations. An external tracker (*ET*) is used to track camera (*C*) with the attached target ( $C_{target}$ ). A second target ( $HIP_{target}$ ) is rigidly attached to the stylus of the haptic device (*HD*). The spatial relation  $EH$  between the coordinate system  $ET_{origin}$  and  $HD_{origin}$  is static.

automatic calibration of the device in a fixed physical location. However, PHANToM Desktop and Premium devices have no inkwell for calibration, which results in unpredictable reset positions. The reported position from the forward kinematic model contains non-linear errors, if the angles between the joints and between the first joint and the  $y$ -axis of the haptic workspace are not equal to 90 degree during reset.

There exists a cyclic dependency between the estimation of the transform  $EH$  and the joint-angle calibration, since the method used for determining  $EH$  requires position measurements from the haptic device. The initial estimation of  $EH$  ( $EH_{init}$ ) contains errors, which are caused by the non-linear errors in the HIP positions. As part of the position calibration, we present a novel step to estimate  $EH_{init}$  without any dependency on the haptic pose (Section 2.2.1).

Each step during the workspace calibration has different requirements on the motion contained in the input data, as well as constraints like a fixed HIP position. Some of those requirements are compatible. In our previous system [12] the users had to separately provide input for each step, which caused unnecessary repetition. In this article, we improve the usability by reducing the required number of inputs to three: gimbal-angles data, reference orientation data, and joint-angles data. The order of data collection is chosen on purpose, since the first two data sets require a fixed position of the  $HIP_{pose}$ .

The haptic device measurements are received with a frame rate of 500-1,000 Hz. In order to ensure precise temporal correspondences, we re-sample the haptic data to match the frame rate of the external tracker (60-100 Hz) using linear interpolation. The resulting stream of correspondences typically consists of several thousand data points. We limit the number of correspondences to  $N = 200$ , in order to avoid over-fitting during the optimization processes. The selection of samples for the optimization processes is an important part for each calibration step and discussed in the relevant sections.

An overview on the workspace calibration procedure is given in Fig. 2. The procedure starts by determining the time-delay between *ET* and *HD* (Section 2.1). The transform  $EH$  and the correction factors for the joint-angles  $\phi_{ja}$  are calculated during the position calibration (Section 2.2). The

orientation calibration provides the correction factors for the gimbal-angles  $\phi_{ga}$  (Section 2.3).

## 2.1 Time Delay Estimation

The external tracker and the haptic device are not synchronized in our system and deliver their measurements with an unknown delay. The purpose of this calibration step is to estimate the delay between *ET* and *HD* in order to compensate it. Measurements from the haptic device are then temporally aligned with the external tracker using linear interpolation.

The basic idea in [17] is to gather measurements with accurate time-stamps from two different sensors that have been transformed into a common coordinate system. Using segments of these measurement streams, the similarity between two signals is computed, while one of the slices is shifted in time. The shift in time, which maximizes the similarity between the signals, corresponds to the time-delay between the sensors.

The time-delay estimation algorithm expects two unsynchronized streams of corresponding 3D positions in a common coordinate system,  $HD_{origin}$  in our case. Therefore we need to estimate the transform between the external tracker  $ET_{origin}$  and haptic device  $HD_{origin}$ . This is done in two steps: First, we estimate the  $HIP_{pose}$  position in tracker coordinates  $HIP_{target}$  using the tooltip calibration method presented by Tuceryan et al. [19]. We prefer a mechanical stand to fix the gimbal position over a fixation force as proposed by Harders for increased accuracy of the reference position. Second, we determine a 6DOF transform between  $ET_{origin}$  and  $HD_{origin}$  using the absolute orientation algorithm presented by Horn [20].

In the next step, the stylus is moved on an arbitrary path with non-constant distance to the workspace origin. The corresponding position measurements from  $HIP_{target}$  and the haptic stylus are then reduced in dimensionality using the euclidean norm (in this case the distance from the workspace origin). The resulting signals are then correlated. In [17] they note that the registration of the two tracking systems does not have to be very precise, since the time-delay estimation is rather robust against spatial registration errors. Thus we can use the same registration procedure used in subsequent steps of the calibration, where the time-delay is not yet compensated. In our setup, the time-delay estimation results in a delay of 9 ms for the PHANToM Omni setup (the measurements of the Naturalpoint OptiTrack system arrive 9 ms after the measurements from the PHANToM Omni on the host computer) and 19ms for the PHANToM Premium 1.5 setup.

Different strategies like buffering, linear interpolation, or Kalman filters can be applied to temporally align the measurements. Since the update rate of PHANToM devices is very high (1 KHz) compared to the reference tracking systems (60-100 Hz), we use a linear interpolation mechanism on the PHANToM measurements to provide corresponding samples for measurements from the reference tracking system. All subsequent steps will now receive temporally aligned measurements.

We are estimating the time delay between the sensors with a precision of 1 ms. The average movement speed during the calibration procedure was  $147 \pm 14 \frac{\text{mm}}{\text{s}}$ . This results

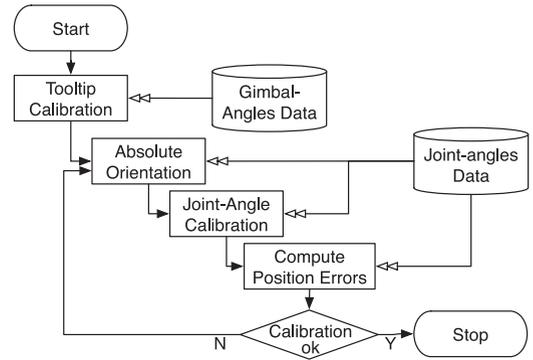


Fig. 4. Position calibration procedure overview. The  $HIP_{pose}$  reference in *HD* coordinates is calculated using the results from the tooltip calibration and the absolute orientation calibration. The reference is then used in an iterative optimization process to minimize the errors in the joint-angle sensors.

in a maximal base error of 0.147 mm from the time delay between the sensors.

## 2.2 Position Calibration

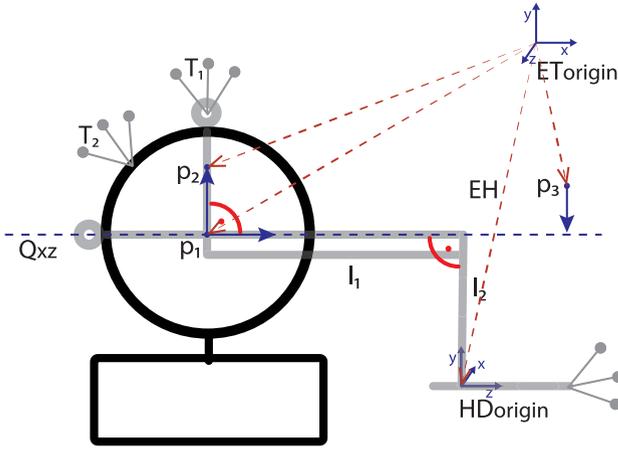
The procedure for compensating the *HIP* position errors is shown in Fig. 4 and is based on the iterative position calibration method presented by Harders [10]. Initially, we determine the *HIP* position relative to  $HIP_{target}$  using the pivot calibration algorithm by Tuceryan et al. [19]. Then, we repeat the estimation of *EH* and the joint-angle calibration until the change in the residual is below a threshold  $\epsilon = 10^{-4}$  mm.

The required inputs for the iterative optimization step are:

- *HIP* reference positions provided via the  $HIP_{target}$  in *HD* coordinates
- corresponding joint-angle measurements from the haptic device
- the haptic device parameters: joint-lengths and origin offset.

Users are advised to collect point correspondences that cover the complete workspace during the data collection. Operators can freely move the haptic stylus to sample data since the relative latency between *ET* and *HD* is already compensated (Section 2.1).

The raw sensor streams are re-sampled and filtered using spatial and temporal constraints to reduce the number of samples. The remaining measurements are clustered into  $N = 200$  clusters to extract an evenly distributed selection. We determine  $N$  centroids from the *HIP* position measurements using the K-Means algorithm with euclidean distance function [21]. The centroids are initialized by randomly picking  $N$  measurements from the original data sets. Therefore, the resulting centroids differ slightly whenever the clustering is performed. For each centroid, a time-stamp with the closest available *HIP* position is then looked up in a KDTree [22]. Finally, the correspondences for the selected time-stamps are retrieved from the re-sampled data streams and used for calibration. The clustering is repeated for each optimization step and yields a varying selection of samples, which makes the calibration more robust against noisy tracking data.



(a) The landmarks  $p_1$ ,  $p_2$ , and  $p_3$  of the haptic device workspace are measured using the external tracker to determine the initial estimation of  $EH$ .



(b) PHANToM Premium with tracking targets attached for reference position estimation.

Fig. 5. Additional targets are rigidly attached to the haptic device for improved reference position estimation. Thus, the initial value of  $EH$  is determined without relying on a correct initialization pose for haptic devices that do not provide a calibration inkwell.

In order to provide the reference position in  $HD$  coordinates, we need a precise estimation of the transform  $EH$  which is discussed in Section 2.2.1. The errors in the joint-angle sensors are then compensated using Harders' method [10] (Section 2.2.2).

### 2.2.1 Reference Position

A crucial step during the position calibration is the computation of the transform  $EH$  between the external tracker  $ET_{origin}$  and the haptic device  $HD_{origin}$ . A commonly used method for computing this transform is the absolute orientation method by Horn [20]. The algorithm requires a set of corresponding 3D positions from both coordinate frames as input and results in the affine transformation  $EH$ , which is required to express measurements from  $ET$  in  $HD$  or vice versa. This transform is used in all subsequent steps to calculate the reference measurements for calibration. High accuracy is therefore essential for the overall calibration result.

As previously discussed, the  $HIP_{pose}$  from the haptic device measurements typically contain non-linear errors, which are caused by the reset procedure and sensor inaccuracies. In the current state-of-the-art calibration, these distorted measurements are used as input for the absolute orientation algorithm. The resulting transform therefore contains errors. This cyclic dependency between  $EH$  and the correction factors for the haptic device is the main reason for the iterative optimization process proposed by Harders et al. [10].

During the development and presentation of our ISMAR 2014 demonstration [23], we collected 65 complete data sets for workspace calibration. We evaluated all data sets with the same parameters (see Section 3) and found that only 55 percent of the sessions resulted in a mean error  $< 2.0$  mm, 22 percent could not be calibrated at all and 23 percent resulted in a mean error  $< 10$  mm. In most sessions, where the calibration fails completely, the causes are likely to be incorrect user input or external influences, like invalid tracking data or bad tracker calibration. The large number of sessions, where position errors are close to an acceptable result, motivated us to revisit the position calibration. We

think that a major problem with Harders' method [10] is the dependency on the uncalibrated position of the haptic stylus ( $HIP_{pose}$ ) for the initial estimation of  $EH$  ( $EH_{init}$ ).

We present a novel algorithm to determine the transform between  $ET_{origin}$  and  $HD_{origin}$ , which does not depend on the  $HIP_{pose}$  of the haptic device. The resulting transform can be used as initial estimation for  $EH$  when calibrating the joint-angles. A better initial guess of  $EH$  reduces the number or required iterations and improves the overall calibration result (see Section 3).

Our algorithm calculates translation  $T_{EH}$  and rotation  $R_{EH}$  of  $EH$  separately based on three landmarks. These landmarks are measured using the external tracker as shown in Fig. 5a:  $p_1$  locates the central joint inside the turret,  $p_2$  defines the direction of the  $y$ -axis in  $HD_{origin}$ , and  $p_3$  is a point in the haptic workspace, which we use to define the  $z$ -axis of  $HD_{origin}$ . Once the  $y$ -axis and  $z$ -axis are known, we can calculate  $R_{EH}$ .  $T_{EH}$  is computed by concatenating  $p_1$  and the translation of the reset pose expressed in the  $ET$  coordinate system.

The points  $p_1$  and  $p_2$  are measured using two additional tracking targets ( $T_1$ ,  $T_2$ ), which are rigidly mounted on the haptic device during the position calibration as shown in Fig. 5b. We are using the tooltip calibration method by Tuceryan et al. [19] to determine  $p_1$  from the spherical movement of  $T_1$  around the central joint and  $p_2$  from the circular movement of  $T_2$  around the  $y$ -axis.  $p_1$  represents the origin of the forward kinematic model and the vector  $\hat{a}_y = \frac{p_2 - p_1}{\|p_2 - p_1\|}$  represents the unit-vector on the  $y$ -axis of the device. Let  $Q_{xz}$  be a plane, that is defined by point  $p_1$  and its normal  $\hat{a}_y$ .

Next, we look up a position  $p_3$  from a  $HIP_{pose}$  measured by the external tracker, where the haptic device reports an angle  $\theta_1 \approx 0$  and project it onto the plane  $Q_{xz}$  to determine the remaining  $z$ -axis and  $x$ -axis unit-vectors:

$$p'_3 = p_3 - ((p_3 - p_1) \cdot \hat{a}_y) \hat{a}_y, \quad (1)$$

$$\hat{a}_z = \frac{p'_3 - p_1}{\|p'_3 - p_1\|}, \quad (2)$$

$$\hat{a}_x = \hat{a}_y \times \hat{a}_z. \quad (3)$$

The rotation matrix  $R_{EH}$  is then directly composed from the axis unit-vectors:

$$R_{EH} = [\hat{a}_x \hat{a}_y \hat{a}_z]. \quad (4)$$

The accuracy of the selection of  $p_3$  is not critical, since the forward kinematics model does not make any assumptions on the direction of  $\hat{a}_z$ , as long as it is contained in  $Q_{xz}$ .

Once the axes  $\hat{a}_y$  and  $\hat{a}_z$  are known, we can calculate  $T_{EH}$ :

$$T_{EH} = p_1 + l_1 \hat{a}_z - l_2 \hat{a}_y. \quad (5)$$

$EH_{init}$ , the initial estimation of  $EH$ , is then given by:

$$EH_{init} = \begin{bmatrix} & R_{EH} & & T_{EH} \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (6)$$

We have compared the position and orientation errors for the complete calibration process with and without the initial estimation of  $EH$ . In most cases, the new initialization method results in improved calibration quality, unless the initial error is already very small (see Section 3).

### 2.2.2 Joint-Angles Error Compensation

Once the initial transform  $EH$  is known, the iterative optimization of the joint-angles can be performed. In each iteration, the correction factors ( $\phi_{ja}$ ) for the joint-angles are computed using the Levenberg-Marquardt algorithm to minimize the remaining position error. With the resulting joint-angle correction factors, a new data set is computed, clustered, and used as input for the refinement of  $EH$  with higher precision. The optimization process continues until the change of the mean position errors between iterations is below  $10^{-4}$  mm or the maximum number of iterations is reached. The process typically terminates after four iterations with our new method for determining  $EH$  and after 10 without.

The measurement stream for  $EH$  estimation is filtered using the following criteria: distance  $d_o$  of  $HIP_{pose}$  measurements to  $HD_{origin}$  is  $\leq 10$  cm, minimum euclidean distance between samples of  $HIP_{pose}$   $d_d > 5$  mm. For each iteration, we relax the constraint  $d_d$  slightly in order to vary the sample selection:  $d_d = 0.9d_d$ .

The input for joint-angles calibration is filtered in a similar way, but instead of the  $d_o$  constraint, we remove outliers in the reference position by limiting the distance  $d_r$  between the haptic pose and the reference.

During the iterative optimization, we found that subsequent joint-angle calibrations could result in large errors, from which the optimization process cannot recover. This can either be caused by biased sample selection, tracker noise, or local minima in the optimization space. We retry the optimization with the previous calibration results and a re-clustered sample set until three consecutive optimizations fail to improve, in order to recover from invalid optimization results.

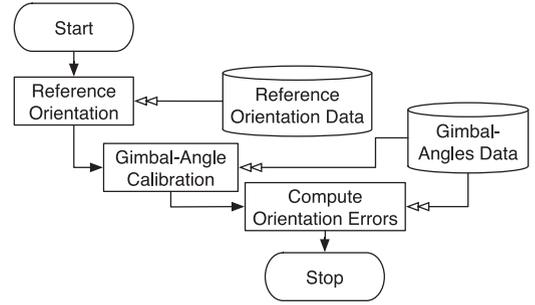


Fig. 6. Orientation calibration procedure overview. A reference orientation is determined with the proposed method and then used to minimize the gimbal-angle sensors errors.

To this end, we have reduced the average position error significantly. The resulting transform  $ET$  and the correction factors for joint-angles are then used for orientation calibration.

### 2.3 Orientation Calibration

The calibration of gimbal angles requires a known reference that can be used to determine the correction parameters for the sensor readings by minimizing the orientation error between the reference and the calculated stylus pose. Such a reference that represents the physical orientation of the haptic stylus can be determined in two ways: using a custom, precisely manufactured, and calibrated tracking target with well-known properties, or by exploiting the mechanical and geometrical properties of the haptic stylus with attached tracking target. The first approach is expensive and not easily adaptable to new combinations of tracking systems and haptic devices. Therefore, we chose the latter approach, which is more general. In this section, we present the algorithm to determine the reference orientation, then we show how the reference orientation can be used to calibrate the gimbal angles. An overview on the process is given in Fig. 6.

Precise temporal alignment of the pose correspondences from the haptic device and the external tracker is essential for our approach. By comparing the poses reported from the haptic device with the physical pose in a high-speed video (120 Hz), we found, that the gimbal sensor readings are delayed. Geomagic support acknowledged this undocumented behavior and stated that the gimbal measurements are by default interpolated over 180 frames to avoid jitter. This interpolation leads to delayed changes in orientation and therefore incorrect pose measurements. The interpolation can be configured only by editing the parameter `FilterGimbalSize` in the device configuration file. We have set this parameter to a value of 0 to avoid any delay in the reported orientation. The parameter change did not introduce any noticeable jitter on our device.

Furthermore, the forward kinematics equations from Çavuşoğlu et al. [11] need to be extended to include the gimbal angles in order to compute a full 6DOF  $HIP_{pose}$  (see Fig. 7). We report all equations in right-handed coordinates. Shortcuts are provided for  $\sin(\cdot) = s(\cdot)$  and  $\cos(\cdot) = c(\cdot)$ .

Let  $l = (l_1, l_2)$  be the vector of joint lengths and  $\theta_{ja} = (\theta_1, \theta_2, \theta_3)$  the vector of joint angles of the haptic device. The translation from the haptic workspace origin to the HIP position  $T_{HIP}$  is given by:

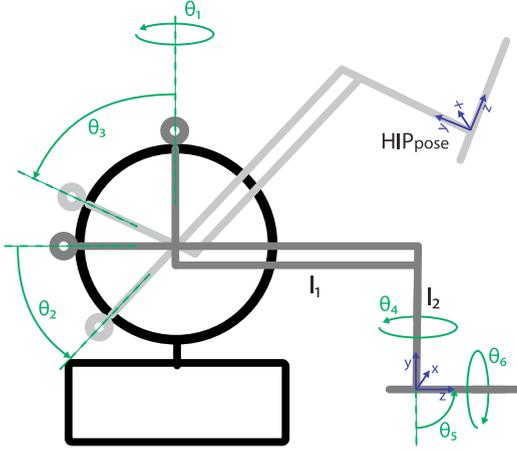


Fig. 7. Schema of the mechanical parts of a PHANTOM haptic device for deriving the 6DOF forward kinematic model.

$$T_{HIP}(\theta_{ja}, l) = \begin{bmatrix} -s(\theta_1)(c(\theta_2)l_1 + l_2s(\theta_3)) \\ -c(\theta_3)l_2 + l_1s(\theta_2) + l_2 \\ c(\theta_1)(c(\theta_2)l_1 + l_2s(\theta_3)) - l_1 \end{bmatrix} \quad (7)$$

and the orientation of the second arm by:

$$R_{123}(\theta_{ja}) = \begin{bmatrix} c(\theta_1) & s(\theta_1)s(\theta_3) & -s(\theta_1)c(\theta_3) \\ 0 & c(\theta_3) & s(\theta_3) \\ s(\theta_1) & -s(\theta_3)c(\theta_1) & c(\theta_1)c(\theta_3) \end{bmatrix}. \quad (8)$$

As detailed in Fig. 7 the rotations  $\theta_{ga} = (\theta_4, \theta_5, \theta_6)$  measured by the gimbal sensors can be expressed as:

$$R_4 = \left( \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \theta_4 \right) R_5 = \left( \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, -\theta_5 \right) R_6 = \left( \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, -\theta_6 \right), \quad (9)$$

where  $(v, \theta)$  refers to a rotation around axis  $v$  with an angle  $\theta$ . The rotations are concatenated, which results in a complete rotation  $R_{HIP}$ :

$$R_{HIP}(\theta_{ja}, \theta_{ga}) = R_{123}R_4R_5R_6 \quad (10)$$

and a full 6DOF pose for the haptic stylus  $HIP_{pose}$ :

$$HIP_{pose}(\theta_{ja}, \theta_{ga}, l) = \begin{bmatrix} R_{HIP} & T_{HIP} \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (11)$$

### 2.3.1 Reference Orientation

The physical orientation of the haptic stylus is initially unknown. We present a novel approach to determine the stylus orientation in relation to  $HIP_{target}$ .

When the stylus is rotated around its longitudinal axis, the attached tracking target and its markers travel on a circular path around the z-axis (see Fig. 8a). The centers of these circles define a straight line  $REF_{zaxis}$  through the HIP and represent the physical orientation of the stylus. Aligning the z-axis of  $HIP_{pose}$  with  $REF_{zaxis}$  during the optimization yields correction factors for  $\theta_4$  and  $\theta_5$ .

Another property of the haptic stylus is that the rotation around its longitudinal axis is mechanically limited in both directions. We use this property to calibrate  $\theta_6$ .

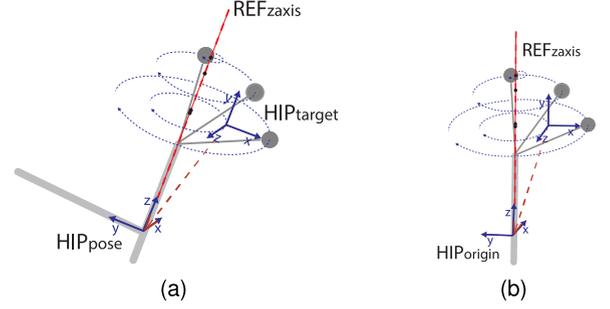


Fig. 8.  $HIP_{target}$  origin and its markers travel on a circular path around the physical z-axis of the haptic stylus when it is rotated around its longitudinal axis. The centers of the circles, which are marked by black dots, are on a straight line through the  $HIP_{pose}$  position. (a) stylus rotation in tracker coordinates. (b) stylus rotation in  $HIP_{origin}$  coordinates.

When users rotate the stylus around  $REF_{zaxis}$  without any mechanical fixture, they cannot completely avoid a change in orientation on the other axes. Therefore, we cannot simply fit the circles on raw measurements. Since we have corresponding sensor readings from joint and gimbal sensors of the haptic device, we can compensate this movement using the forward kinematic pose of  $\theta_1$  to  $\theta_5$

$$HIP'(\theta_{ja}, \theta_{ga}, l) = \begin{bmatrix} R_{123}R_4R_5 & T_{HIP} \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (12)$$

Let  $n$  be the number of correspondences in this step,  $v = (v_i, i = 1 \dots n)$  be the vector of  $HIP_{target}$  marker positions,  $P = (P_i, i = 1 \dots n)$  a vector of  $HIP_{target}$  poses,  $F = (F_i, i = 1 \dots n)$  a vector of the corresponding poses calculated using  $HIP'$ , and  $EH^{-1}$  the transform between  $ET_{origin}$  and  $HD_{origin}$ . As shown in Fig. 8b, each  $v_i$  is transformed into stylus coordinates  $HIP_{origin}$  using

$$v'_i = F_i^{-1}(EH^{-1}v_i). \quad (13)$$

As shown in Fig. 9, circles can now be fitted in the x-y plane using orthogonal distance regression [24]. This results in a series of 3D points defining the z-axis  $z$  in stylus coordinates, which is calculated using singular value decomposition. To obtain  $z$  in  $ET$  coordinates ( $z_{ref}$ ), we transform  $z$  using  $n$  correspondences of  $P$  and  $F$

$$z_{ref} = \frac{1}{n} \sum_{i=1}^n (EH^{-1}P_i)^{-1}(F_i z). \quad (14)$$

At this stage, we obtained one reference axis, which allows us to minimize errors for  $\theta_4$  and  $\theta_5$  (see Section 2.3.2). Next, we determine the reference needed for calibrating  $\theta_6$ .

The stylus rotation around its z-axis is mechanically limited. We advise users to rotate the stylus towards both extremes. This mechanical limit is also reflected in the

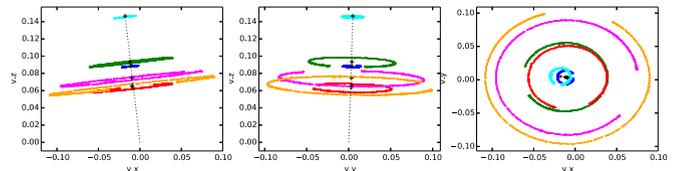


Fig. 9.  $REF_{zaxis}$  is determined by finding centers of circles described by markers and fitting a line through the stylus origin.

circular paths described by the  $HIP_{target}$  markers as shown in Fig. 9. Each path has a gap that corresponds to the mechanical limit of the rotation. The opposite side of the center of the gap defines the origin of  $\theta_6$ , while the upper and lower bounds define its range. We use this correspondence to calibrate  $\theta_6$ .

Let  $C = (C_i, i = 1 \dots n)$  be the vector of 2D positions in the x-y plane of the circle with the largest value for radius  $r$  divided by its residual. We transform  $C$  into a unit circle  $C' = (C'_i, i = 1 \dots n)$  using  $C'_i = \frac{C_i}{r}$  and compute the vector  $\alpha = (\alpha_i, i = 1 \dots n)$  consisting of angles that correspond to each  $C'_i$ . The lower bound  $\alpha_l$  and the upper bound  $\alpha_u$  of the gap in  $C'$  is determined using a histogram with 360 bins. The origin reference for  $\theta_6$  is  $\alpha_o = \frac{\alpha_u - \alpha_l}{2} + \pi$ . As a final step, we compute the reference vector for calibration as  $\hat{\alpha} = -\alpha + \alpha_o$ .

In the next section, we discuss the error minimization for gimbal angles using  $z_{ref}$  and  $\hat{\alpha}$ .

### 2.3.2 Gimbal-Angles Error Compensation

We consider two models for the deviation between gimbal angles and sensor measurements. First, a linear function  $\theta_{ga} = k_{ga}\psi_{ga} + m_{ga}$  where  $\psi_{ga}$  represents the vector of measured gimbal angles,  $k_{ga}$  the gains, and  $m_{ga}$  the offsets. Second, we model the gimbal angle errors using a quadratic function  $\theta_{ga} = j_{ga}\psi_{ga}^2 + k_{ga}\psi_{ga} + m_{ga}$  with an additional factor  $j_{ga}$  for the squared angle. We evaluated both error models, because gimbal sensors in PHANToM Omni and Desktop devices are potentiometers with a linearity of  $\pm 5$  percent as opposed to digital encoders for the joint angles.

Let  $\phi_{ja} = (k_{ja}, m_{ja})$  be the correction factors for  $\psi_{ja}$ , which have been determined during the position calibration. Let  $\phi_{ga} = (j_{ga}, k_{ga}, m_{ga})$  be the unknown correction factors for  $\psi_{ga}$ , setting  $j_{ga} = 0$  for the linear error model. The corrected stylus rotation can be expressed as  $R'(\psi_{ja}, \phi_{ja}, \psi_{ga}, \phi_{ga})$  by substituting  $\phi_{ja}$  and  $\phi_{ga}$  into (10).

First we obtain the correction factors  $\phi_{45}$  for  $\theta_4$  and  $\theta_5$  by performing a nonlinear optimization using the Levenberg-Marquardt algorithm. This minimizes the angle between the z-axis of  $HIP_{pose}$  and  $\hat{z} = R_{target} \cdot z_{ref}$ .

$$\min_{\phi_{45}} \sum_{i=1}^n 1 - \hat{z}_i \cdot \left( R'_i(\psi_{ja}, \phi_{ja}, \psi_{ga}, \phi_{45}) \cdot \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right). \quad (15)$$

Determining the correction factors  $\phi_6$  for  $\theta_6$  is straightforward, since we have obtained direct correspondences in  $\hat{\alpha}$ . Let vector  $\beta$  be the vector of  $\theta_6$  measurements. We minimize the error using least-squares optimization

$$\min_{\phi_6} \sum_{i=1}^n (j_6\beta_i^2 + k_6\beta_i + m_6) - \hat{\alpha}_i. \quad (16)$$

The final correction factors  $\phi_{ga}$  are then composed from  $\phi_{45}$  and  $\phi_6$ .

## 3 EVALUATION

In this section we present the numerical evaluation of repeated calibrations for our setups and discuss the results. We collected five data sets to validate various aspects of our workspace calibration method:

- **DS<sub>1</sub>**: The first data set was recorded during the evaluation<sup>1</sup> of [12] with measurements from the low-fidelity setup. It is used to evaluate the position accuracy and the orientation calibration method.
- **DS<sub>2</sub>**: The second data set was also recorded during the evaluation<sup>1</sup> of [12] with measurements from the high-fidelity setup. It is used to evaluate the orientation calibration method.
- **DS<sub>3</sub>**: For the third data set, a large number of calibration sessions were recorded during the development and presentation of our ISMAR 2014 demonstration [23] using the *HF* setup.
- **DS<sub>4</sub>**: The fourth data set was recorded with the *HF* setup in order to show the benefits of our new initialization method for the position calibration.
- **DS<sub>5</sub>**: In the fifth data set, we capture sensor measurements in the *HF* setup using a mechanical tracker together with the external tracking system. This allows us to explore the limits of our calibration method.

The *LF* setup covers a small workspace and consists of a PHANToM Omni, a Naturalpoint OptiTrack IR tracking system, and a PointGrey Flea2 firewire camera using a single workstation (Intel Core I7 3.5 GHz, 16 GB RAM, NVidia GForce GTX 570). The *HF* setup is a room-size installation with a PHANToM Premium 1.5 6DOF, an A.R.T ART-TRACK IR tracking system connected via 1 GB/s network and a IDS uEye usb camera on a workstation (Intel Core I5 3.4 GHz, 8 GB RAM, NVidia GForce GTX 460).

We implemented our calibration procedure using the tracking and sensor fusion framework Ubitrack [25], which implements an asynchronous dataflow network and provides components for calibration, persistence, and device drivers. Ubitrack also facilitates modeling of concatenated transforms, sensor fusion, interpolation, and filtering.

During all calibration sessions, we recorded the raw sensor data streams for each calibration step. The sensor streams were stored with precise, latency-compensated timestamps. This enables us to re-perform the calibration process off-line in order to verify and improve our procedure. All results presented in this section are computed off-line with the latest implementation of our calibration procedure as presented in this article.

The errors in  $DS_1$ ,  $DS_2$  and  $DS_3$  are calculated using the original data streams recorded during the calibration process excluding  $N$  samples, which were used during the last optimization step.  $DS_4$  additionally contains a separate stream that was collected after calibration. It is used for numerical evaluation.

The largest collected data set is  $DS_3$  with 65 complete calibration sessions from the *HF* setup. All sessions contain enough data to perform a complete calibration procedure, except for the initial estimation of *EH* as proposed in Section 2.2.1.  $DS_3$  contains sessions with invalid data. The initialization of the haptic device was performed with varying deviations from the optimal reset pose.

Based on the resulting mean position error  $pe_{cal}$  of each session, we split the data set into three groups: good,

1. We re-evaluated all sessions using the original raw sensor data streams and our latest calibration pipeline, that contains a complete implementation of Harders' [10] iterative position calibration method.

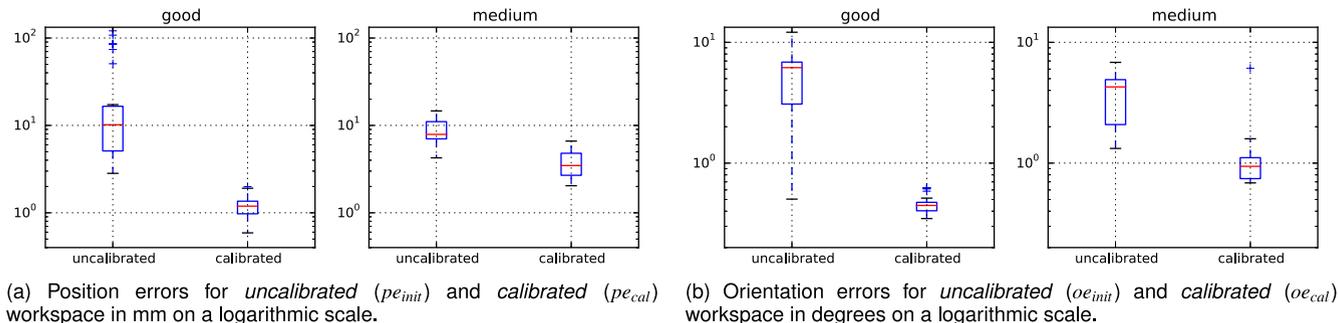


Fig. 10. Evaluation of  $DS_3$  using the  $HF$  setup: sessions are grouped by resulting position error (good  $< 2$  mm  $\leq$  medium  $< 10$  mm).

medium, and bad. With our current implementation, we can successfully calibrate 55 percent of the sessions resulting in a mean position error  $pe_{cal} < 2.0$  mm, 23 percent result in  $pe_{cal} < 10$  mm, and 22 percent cannot be calibrated at all. The results for position errors are shown in Fig. 10a and the results for orientation errors are shown in Fig. 10b. The numerical results are given in Table 1.

The results show that the quality of the position calibration has an effect on the orientation calibration, which was expected because of the forward kinematic pose calculation. More interestingly the diagrams show that most sessions with medium position calibration results have initial errors similar to the good sessions. We believe that the main reason for achieving only medium quality results in this group is the suboptimal initialization of the haptic device since the workspace setup itself remained mostly unchanged.

In Section 3.1 we present the results of the position calibration using  $DS_1$  and the evaluation of  $DS_4$ , which shows improved position calibration results with our new method for estimating  $EH_{init}$  in comparison to Harders' original method. In Section 3.2, we discuss the results of orientation calibration using  $DS_1$ ,  $DS_2$  and  $DS_4$ . In Section 3.3, we highlight improvements to the usability of our workspace calibration method. Section 3.4 concludes our evaluation with the comparison of calibration results in  $DS_5$  from the external tracker with the results from a highly precise mechanical tracker.

### 3.1 Joint-Angle Calibration Results

We report the position error as euclidean distance between the reference HIP position and the forward kinematic pose from the haptic stylus. For the  $LF$  setup, we verified the position calibration using  $DS_1$ , which contains 12 sessions.

TABLE 1  
Statistical Results for the Grouped Data Set  $DS_3$

	good	medium	bad
count	36	15	14
$pe_{init}$	$23.8 \pm 33.2$ mm	$8.9 \pm 3.0$ mm	$207.6 \pm 65.1$ mm
$pe_{cal}$	$1.2 \pm 0.4$ mm	$3.8 \pm 1.5$ mm	$95.8 \pm 38.6$ mm
$oe_{init}$	$5.1 \pm 2.5^\circ$	$3.7 \pm 1.8^\circ$	$24.8 \pm 15.6^\circ$
$oe_{calib}$	$0.5 \pm 0.1^\circ$	$1.3 \pm 1.4^\circ$	$21.1 \pm 16.6^\circ$
$te_{et}$	$0.98 \pm 0.26$ mm	$2.23 \pm 3.60$ mm	$29.11 \pm 42.13$ mm

$pe_{init}$ : position error in mm before calibration,  $pe_{cal}$ : position error in mm after calibration,  $oe_{init}$ : orientation error in degrees before calibration,  $oe_{calib}$ : orientation error in degrees after calibration,  $te_{et}$ : mean back-projection error in mm of tooltip transform during pivot calibration.

The uncalibrated error  $\overline{pe_{init}} = 5.23 \pm 0.87$  mm was reduced to  $\overline{pe_{cal}} = 1.51 \pm 0.476$  mm. All position error evaluations for  $HF$  were conducted with  $DS_3$  as shown in Fig. 10a and  $DS_4$ , which contains the required data to calculate  $EH_{init}$  as described in Section 2.2.1.

For the data collection of  $DS_4$  we mounted additional tracking targets ( $T_1, T_2$ ) to the haptic device during position calibration. We collected data from 18 sessions and performed the iterative optimization for the joint-angles with ( $fwkbase$ ) and without ( $harders$ ) the new estimation of  $EH_{init}$ . The results are shown in Fig. 11a.

We applied the Kruskal-Wallis test to check whether there is a significant difference in mean position errors between  $harders$  and  $fwkbase$  for each session. We found that in all cases but one,  $fwkbase$  significantly differs from  $harders$  with  $p < 0.05$  and that  $fwkbase$  resulted in better position calibration in 14 sessions. The average position error for  $uncalibrated$  is  $54.52 \pm 96.37$  mm, improved in  $harders$  to  $1.73 \pm 1.46$  mm, and most accurate in  $fwkbase$  with  $1.07 \pm 0.13$  mm. The large error for the  $uncalibrated$  case is caused by the fact that we deliberately used extreme device configurations during haptic device initialization to test the estimation of  $EH_{init}$ , but calculated the position error using the final, most accurate result of  $EH$ . The mean RMS error for tooltip calibration in  $DS_4$  is  $1.04 \pm 0.09$  mm.

Compared to our previous work [12], we could further reduce the remaining mean position error after calibration by 16 percent for the  $LF$  setup and by 58 percent for the  $HF$  setup.

### 3.2 Gimbal-Angle Calibration Results

We measure the orientation error as angle between the z-axis of the reference pose and the haptic pose. This metric ignores errors in the rotation around the longitudinal axis of

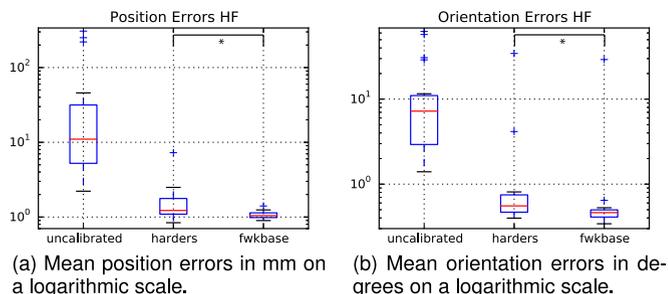


Fig. 11. Results for the  $HF$  setup in  $DS_4$ : *Uncalibrated* errors are compared to the iterative optimization process without (*harders*) and with (*fwkbase*) the new method for estimating  $EH_{init}$ . Horizontal bars indicate significant effects: \* denotes  $p < .05$ .

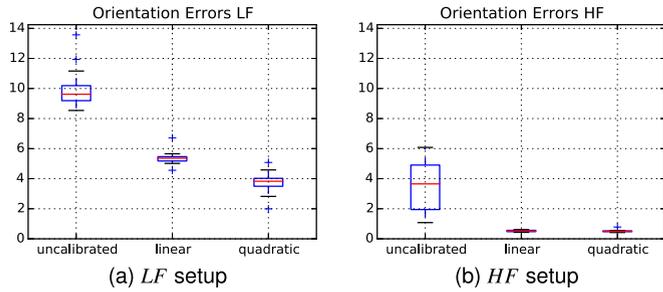


Fig. 12. Results for the data sets  $DS_1$  and  $DS_2$ : Mean orientation errors in degrees for three conditions: *uncalibrated*, gimbal angle calibration with *linear* error model, gimbal angle calibration with *quadratic* error model

the haptic stylus, but closely resembles the requirement for accurate overlays. For the evaluation of orientation errors we analyze the data sets  $DS_1$  with twelve sessions from the *LF* setup, and  $DS_2$  with six sessions from the *HF* setup. These data sets were originally recorded for our ISMAR 2014 paper [12] and re-evaluated with a corrected implementation of the iterative position calibration algorithm.

Errors are calculated using the data-stream for gimbal-angle calibration excluding all samples that were used during the last optimization. The haptic pose is calculated using the joint-angles and gimbal-angles correction factors. The reference pose is calculated using the result for *EH* from the position calibration step.

We report the orientation error for  $\theta_4$  and  $\theta_5$  in degrees between the estimated  $REF_{zaxis}$  and the z-axis of the calculated  $HIP_{pose}$  for three conditions in both setups: *uncalibrated*, *linear*: gimbal angle calibration with linear error model, and *quadratic*: gimbal angle calibration with quadratic error model. The resulting orientation errors are shown in Fig. 12.

In  $DS_1$  the average orientation error for *uncalibrated* is  $10.03 \pm 1.49^\circ$ , improved in *linear* with  $5.38 \pm 0.51^\circ$ , and most accurate in *quadratic* with  $3.74 \pm 0.79^\circ$ . In  $DS_2$  the average orientation error for *uncalibrated* is  $3.53 \pm 2.0^\circ$ , and significantly improved in *linear* to  $0.52 \pm 0.07^\circ$  and in *quadratic* with  $0.54 \pm 0.13^\circ$ .

Fig. 13 visualizes the orientation calibration results within the workspace. Four heat-maps show the reduction of orientation errors in degrees at four positions of the haptic workspace using an unwrapped spherical heat-map with hammer projection; higher values represent more error compensation.

Fig. 14 illustrates the distribution of the orientation error caused by  $\theta_4$  and  $\theta_5$  for a typical calibration session in *LF* and *HF*. Measurements were recorded with a fixed stylus position close to the workspace origin. The error in spherical coordinates ( $HD_{origin}$ ) is visualized using a heat-map that is projected using a hammer projection. White areas contain no measurements due to mechanical limits of the haptic stylus; higher values represent a larger error.

The results for gimbal calibration suggest that a linear error model is sufficient for PHANToM Premium models. PHANToM Omni and Desktop devices should be calibrated using the quadratic model. This is in line with our initial expectations. In our previous work [12] we got contradicting results, because we used the quadratic error model for

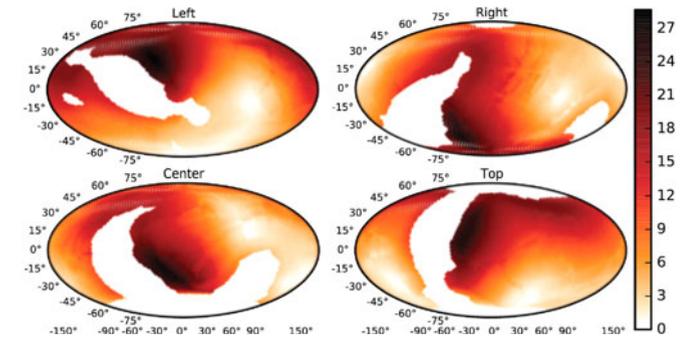


Fig. 13. Visualization of reduced orientation errors at four positions of the haptic workspace (in degrees on an unwrapped spherical heat-map with hammer projection).

position and orientation calibration, which probably led to over-fitting of the joint-angles correction factors.

The results of the orientation calibration for  $DS_4$  are shown in Fig. 11b. We applied the Kruskal-Wallis test to check whether there is a significant difference in mean orientation errors between *harders* and *fwkbase* for each session. We found that in 12 out of 18 cases *fwkbase* significantly differs from *harders* with  $p < 0.05$  and resulted in better orientation calibration. The average orientation error for *uncalibrated* is  $14.26 \pm 18.66^\circ$ , improved in *harders* to  $2.64 \pm 7.98^\circ$ , and most accurate in *fwkbase* with  $2.06 \pm 6.77^\circ$ .

We have shown that the gimbal calibration significantly improves the accuracy of the reported stylus orientation, which enables precise overlays and improves torque feedback co-location. Compared to our previous work [12], we could further reduce the remaining mean orientation error after calibration by 34 percent for the *LF* setup and by 86 percent for the *HF* setup.

### 3.3 Usability Improvements

Our results indicate that time-delay estimation can be used to temporally align sensor input from unsynchronized sources, which streamlines the data collection process and can improve workspace coverage. Our method improves usability of the workspace calibration compared to previous works because it enables users to continuously move the haptic stylus during sample collection, as opposed to the commonly used point and hold method. As a result, we can capture more samples for calibration while at the same time simplifying the user interaction. Our data processing pipeline automatically selects a constant number of samples from the recorded data streams, which are evenly distributed in the area of interest and slightly randomized during each selection to reduce the effects of sensor noise.

An important indicator for the usability of the calibration method is the total time required for sample collection where users manipulate the haptic device according to the instructions of each step in the procedure. Compared to our previous work [12], we have reduced the required number of steps for workspace calibration as follows. First, we separate time-delay estimation from workspace calibration since it is not needed unless a change in the setup occurs. Second, we record realtime data to avoid duplicate user input and select the required data for each calibration during a separate offline calibration step. Finally, our improved calibration method is more robust against sensor noise and

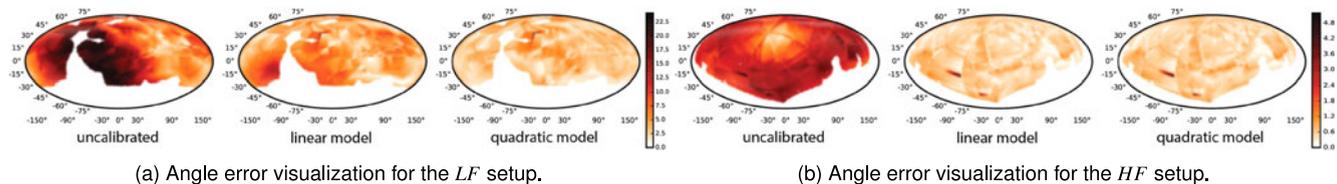


Fig. 14. Angle error in degrees heat-map before and after the calibration measured near the workspace center and displayed in spherical coordinates using hammer projection; white areas contain no measurements.

outliers, which enables faster movement during sample collection.

We compared the duration of user input for our improved calibration method ( $dt_3$ ) with our previous approach ( $dt_1$ ). The duration was calculated as the sum of the durations of each required step using the timestamps contained in the realtime data. Our new method reduces the average time required for calibrating  $HF$  by more than 50 percent from  $dt_1 = 710 \pm 117s$  in  $DS_2$  to  $dt_3 = 312 \pm 31s$  in  $DS_4$ .

### 3.4 Verification with Mechanical Tracking Arm

We have also validated our calibration procedure with an extended HF setup using the high precision Faro fusion mechanical tracking arm, see Fig. 15a. We use this setup, which provides ground truth for reference measurements, to explore the limits of our current setup.

The highly precise Faro Fusion mechanical tracking arm ( $MT$ ) is at least one magnitude more accurate than the optical tracking system. Using  $MT$ , we can rule out line of sight problem of the optical tracking system during the sample collection. Furthermore, the sensor noise is reduced and the update rate is increased to 240 Hz (four times higher than  $ET$ ).

A major challenge in this setup was the construction of an extremely precise and rigid link between  $MT$  and the haptic stylus. We manufactured a custom connector ( $MT_{con}$ ) using a lath and an NC router that perfectly fits the extension fixture on the haptic stylus as well as the probe holding fixture of the Faro arm (see Fig. 15b). We took great care that there is no play between the parts and that the connector sits perfectly centered on the haptic stylus.

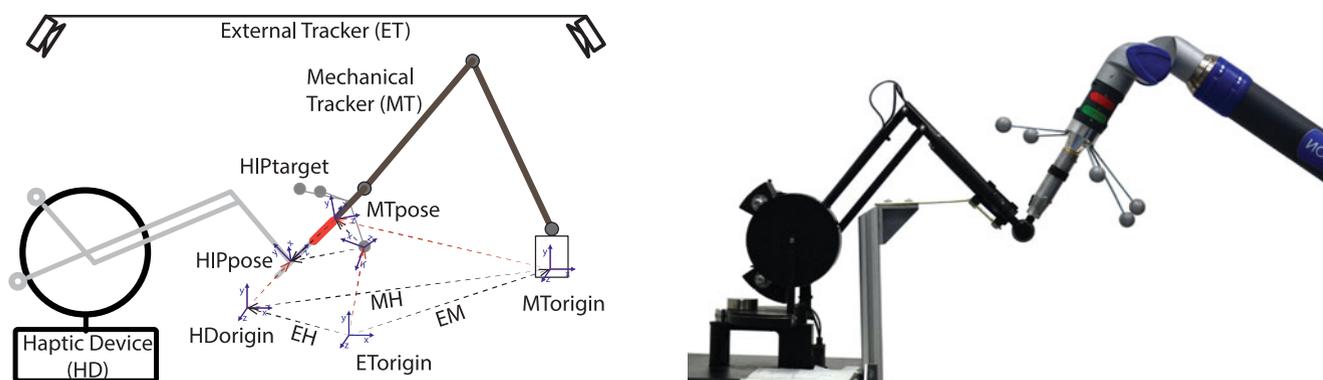
A limitation of  $MT$  is the missing degree of freedom at its tip. Therefore, the haptic stylus cannot be rotated around its

longitudinal axis once it is connected to  $MT$ . This limitation prevents us from performing the calibration step for determining the orientation reference (Section 2.3.1). Instead, we exploit the known mechanical properties of  $MT_{con}$  to provide a precise reference for orientation calibration.

Since we are interested in the comparison of calibration results between  $MT$  and  $ET$ , we also integrated an IR tracking target into  $MT_{con}$ . This setup enables us to record sensor data streams during a calibration session from both reference sources at the same time. We recorded seven calibration sessions using this setup in data set  $DS_5$ .

The calibration procedure closely follows the steps outlined in this article. As a first step, we additionally determine the affine transform ( $S$ ) between the external tracking target pose  $HIP_{target}$  and the pose reported by  $MT$  ( $MT_{pose}$ ) using the hand-eye calibration method from Tsai and Lenz [18]. Then, we estimate the time-delay between the three inputs,  $HD$ ,  $ET$ , and  $MT$ , in order to temporally align them. Next, we separately calibrate the joint-angles with both reference sources separately and determine  $EH$  and the transform  $MH$  between the workspace origin of  $MT$  ( $MT_{origin}$ ) and  $HD_{origin}$  using the iterative position calibration method. We skip the initial estimation of  $EH$  using our new method, since we cannot measure the required landmarks with  $MT$  directly. Instead, we take great care to perform accurate initialization of the haptic device. Finally, we calibrate the gimbal angles using an orientation reference that has to be determined from the mechanical properties of  $MT_{con}$ .

The orientation reference is calculated as follows. First we perform the device calibration of  $MT$  with the default ball-point probe using the vendor supplied tool. Once the calibration is applied,  $MT$  accurately reports the position of



(a) Our extended setup and its spatial relations. We additionally connect the measurement probe of  $MT$  to the haptic stylus.

(b) Connection between mechanical tracker and haptic device with additional IR target.

Fig. 15. VHAR Setup with precise mechanical tracker for results verification. We constructed a mount that connects the haptic stylus to the mechanical tracker and also integrates an IR tracking target. This setup enables us to process temporally aligned real-time sensor data from  $MT$  and  $ET$  in order to compare calibration quality.

the ball-point probe center as sensor output. The orientation of  $MT_{pose}$  is not calibrated with the vendor supplied tool and is therefore unknown. Next, we replace the default ball-point probe with  $MT_{con}$  and perform a tooltip calibration in order to find the HIP position in relation to  $MT_{pose}$ . Due to the construction of  $MT_{con}$ , the reference z-axis is now defined by the vector between the calibrated tip and the center of the ball-point probe. We directly use this reference to calibrate gimbal-angles with  $MT$  and transform it into  $ET$  coordinates using  $S^{-1}$ .

In order to estimate the expected accuracy, we analyzed several aspects of our setup and the data collected in  $DS_5$ . The PHANToM Premium 6DOF has a nominal position resolution of 0.03 mm, the Faro Fusion has a volumetric accuracy of 0.061 mm and repetition error of 0.043 mm. The volumetric accuracy of the ART tracking system depends on the number and position of the IR cameras and the room calibration using the vendor tools. In our setup, judging from the residual errors of the room calibration, we are able to achieve sub-millimeter accuracy in the workspace of the haptic device.

Besides the hardware limitations, there are also limitations with our method, since errors in the temporal alignment and in the calculation of reference measurements accumulate. We have investigated two sources of errors: temporal misalignment after time-delay compensation caused by the movement of the calibration target and tooltip calibration accuracy. Another source of errors are imprecise timestamps that are caused by variable latency in third-party drivers, network, or hardware. Any non-constant delays cannot be compensated with our method.

The spatial accuracy of our time-delay compensation component depends on two factors: the resolution of the delay estimation algorithm, which is 1 ms in our implementation as suggested by Huber et al. [17], and the velocity of the target during calibration. The average speed of the HIP position during joint-angle calibration in data set  $DS_5$  is  $d_t = 0.13 \pm 0.05 \frac{\text{mm}}{\text{ms}}$  with a maximum of  $0.35 \frac{\text{mm}}{\text{ms}}$ . We have also evaluated the velocities of all correspondences used for joint-angle calibration, which results in  $d_t = 0.15 \pm 0.01 \frac{\text{mm}}{\text{ms}}$  with a maximum of  $0.16 \frac{\text{mm}}{\text{ms}}$ .

The tooltip transform between reference pose and the HIP is used to calculate all HIP reference positions. During tooltip calibration, the HIP position is locked using a mechanical fixture to avoid any movement (See Fig. 15b). We validated the accuracy of our setup by applying the tooltip transform to all reference poses recorded during calibration and computed the RMS error of the resulting HIP position for  $MT$ :  $te_{mt} = 0.53 \pm 0.08 \text{ mm}$  and for  $ET$ :  $te_{et} = 0.62 \pm 0.05 \text{ mm}$ . The resulting error  $te_{et}$  in  $DS_5$  is significantly smaller than in  $DS_3$  and  $DS_4$ . We investigated the rigidity of  $HIP_{target}$  that we used for data collection in  $DS_3$  and  $DS_4$  and found that there is a little play between the target and the haptic stylus, which could explain the improved results.

The results for the evaluation of  $DS_5$  are summarized in Table 2. Position errors could be further reduced by more than 35 percent in comparison to  $DS_4$ . The orientation errors are comparable to  $DS_2$  and  $DS_4$ . Although there is no direct comparison possible between  $DS_4$  and  $DS_5$ , the results suggest that our orientation reference algorithm (see Section 2.3.1)

TABLE 2  
Statistical Results of Seven Sessions in  $DS_5$  for the Extended  $HF$  Setup

		$pe$	$oe$
uncalibrated	$ET$	$3.46 \pm 1.20 \text{ mm}$	$5.34 \pm 4.46^\circ$
	$MT$	$4.28 \pm 1.50 \text{ mm}$	$5.29 \pm 4.89^\circ$
iterative	$ET$	$0.77 \pm 0.09 \text{ mm}$	$0.61 \pm 0.05^\circ$
	$MT$	$0.68 \pm 0.08 \text{ mm}$	$0.49 \pm 0.08^\circ$

$pe$ : mean position error in mm,  $oe$ : mean orientation error in degrees.

yields correct results, otherwise the gimbal-angle correction could not be performed in this quality.

There is, however, no significant difference between  $ET$  and  $MT$ , which leads to the assumption that we are still not experiencing the limits of the external tracking system. From our analysis of  $DS_5$  we conclude that extreme care needs to be taken to avoid any mechanical movement during tooltip calibration and at the link between  $HIP_{target}$  and the haptic stylus to achieve optimal calibration results.

We further evaluated the directionality of the remaining position errors in  $DS_5$ . After the calibration, the error vectors point in arbitrary directions. From this we can conclude that our pipeline is working as expected. If time-delay compensation would not work as expected, the error vectors would point into the direction of the movement. An error in the absolute orientation can also be ruled out as the error vectors would point into a common direction, which would be indicated by a positional or rotational offset as shown in [17].

## 4 CONCLUSIONS AND FUTURE WORK

The comprehensive workspace calibration method proposed in this article enables accurate visual overlays on a haptic stylus and precise co-location of haptic feedback in VHAR systems. We presented a novel approach to estimate the physical orientation of a haptic stylus. This enables us to compensate errors in the gimbal sensors, improving orientation accuracy by more than 63 percent. Additionally, the integration of time-delay estimation enables users to calibrate their workspace faster and more conveniently. Our extended position calibration procedure further improves the position accuracy by 38 percent for haptic devices without mechanical fixture for initialization.

The algorithm for gimbal angle calibration and our extensions to the position calibration are generic and can be used for any pen-based haptic device with similar mechanical properties. The integration of time-delay estimation is also applicable to many systems, which integrate multiple unsynchronized sensor inputs.

While in this article we report the remaining position and orientation errors in real world units, the ultimate measure for accurate overlays is the remaining pixel error of augmentations on the users display. The required precision of the haptic device calibration for accurate overlays depends on camera and display parameters, which have not been taken into account in this article. In future work we plan to define a metric to quantify the calibration quality of visuo-haptic augmented reality workspaces similar to [13]. We also plan to study the limits of human perception in regards to co-location error.

The current calibration process still requires a lot of knowledge about the steps and offers no guidance to users. Visual and haptic cues have not been utilized, although they are readily available. In future work, we plan to implement and evaluate haptic and visual cues that guide non-expert users during the calibration process.

Although we compensate the time-delay between the sensors, the temporal alignment errors during fast motion limit the accuracy of our approach. In order to further improve the input to the calibration algorithms, the sample selection procedure could prefer measurements received during slow motion.

Our current implementation does not compensate the time-delay between the camera and the other sensors. Therefore, inconsistent augmentations occur during fast movements of tracked objects or the haptic stylus. The proposed method for estimating the time-delay can also be applied to the camera if fiducial markers are used for viewpoint estimation.

If no external tracker is used at all and the camera viewpoint estimation is solely based on fiducial markers, the calibration could also be performed by attaching a fiducial marker target to the haptic stylus. This would require slow movement during calibration and accurate temporal alignment of the haptic device sensor data and the camera poses.

Finally, improved position and orientation accuracy results in higher fidelity visual and haptic augmentations, which are crucial for fine-motor tasks in areas such as medical training simulators, assembly planning tools, or rapid prototyping applications. Specifically, augmented reality applications where users interact with virtual and real objects at the same time will benefit from accurate haptic device co-location. Examples include real tools attached to haptic devices used to interact with virtual objects, augmented tools to interact with real objects, or composite elements like an augmented tip on physical stylus. The work presented in this article is an important step towards a complete and user friendly workspace calibration, which is essential for enabling real-world VHAR applications.

## ACKNOWLEDGMENTS

The authors greatly acknowledge the support from our colleagues at Technische Universität München: Pascal Fallavollita from NARVIS for providing access to their PHANToM Premium 1.5 6DOF haptic device and Yuta Itoh from FAR for inspiring discussions while developing the reference orientation algorithm. Ulrich Eck is funded by the Award of the Australian Government International Postgraduate Research Scholarship.

## REFERENCES

- [1] S. M. B. I. Botden, S. N. Buzink, M. P. Schijven, and J. J. Jakimowicz, "Augmented versus virtual reality laparoscopic simulation: What is the difference?" *World J. Surgery*, vol. 31, no. 4, pp. 764–772, 2007.
- [2] M. J. Fu, A. D. Hershberger, K. Sano, and M. C. Cavusoglu, "Effect of visuo-haptic co-location on 3D fits' task performance," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, San Francisco, CA, USA, 2011, pp. 3460–3467.
- [3] T. H. Massie and J. K. Salisbury, "The phantom haptic interface: A device for probing virtual objects," in *Proc. Winter Annu. Meet. Symp. Haptic Interfaces Virtual Environ. Teleoperator Syst.*, Chicago, IL, USA, Jan. 1994, pp. 295–300.
- [4] P. Rhenmora, K. Gajananan, P. Haddawy, M. N. Dailey, and S. Suebnukarn, "Augmented reality haptics system for dental surgical skills training," in *Proc. 17th ACM Symp. Virtual Reality Softw. Technol.*, Hong Kong, 2010, pp. 97–98.
- [5] C. Sandor, T. Kuroki, S. Uchiyama, and H. Yamamoto, "Exploring visuo-haptic mixed reality," Tokyo, Japan, Tech. Rep. PRMU 106, 2007.
- [6] M. Csongei, L. Hoang, U. Eck, and C. Sandor, "ClonAR: Rapid redesign of real-world objects," in *Proc. 11th IEEE Int. Symp. Mixed Augmented Reality*, Atlanta, GA, USA, 2012, pp. 277–278.
- [7] F. I. Cosco, C. Garre, F. Bruno, M. Muzzupappa, and M. A. Otaduy, "Augmented touch without visual obtrusion," in *Proc. 8th IEEE Int. Symp. Mixed Augmented Reality*, Orlando, FL, USA, 2009, pp. 99–102.
- [8] J. Vallino and C. Brown, "Haptics in augmented reality," in *Proc. IEEE Int. Conf. Multimedia Comput. Syst.*, 1999, pp. 195–200.
- [9] M. Ikits, C. D. Hansen, and C. R. Johnson, "A comprehensive calibration and registration procedure for the visual haptic workbench," in *Proc. 9th Eur. Workshop Virtual Environ.*, New York, NY, USA, May 2003, pp. 247–254.
- [10] M. Harders, G. Bianchi, B. Knoerlein, and G. Szekely, "Calibration, registration, and synchronization for high precision augmented reality haptics," *IEEE Trans. Vis. Comput. Graph.*, vol. 15, no. 1, pp. 138–149, Jan./Feb. 2009.
- [11] M. C. Çavuşoğlu, D. Feygin, and F. Tendick, "A critical study of the mechanical and electrical properties of the phantom haptic interface and improvements for high performance control," *Presence: Teleoperators Virtual Environ.*, vol. 11, no. 6, pp. 555–568, 2002.
- [12] U. Eck, F. Pankratz, C. Sandor, G. Klinker, and H. Laga, "Comprehensive workspace calibration for visuo-haptic augmented reality," in *Proc. 17th IEEE Int. Symp. Mixed Augmented Reality*, Munich, Bavaria, Germany, Sep. 2014, pp. 123–128.
- [13] D. Wang, Y. Zhang, W. Zhou, H. Zhao, and Z. Chen, "Collocation accuracy of visuo-haptic system: Metrics and calibration," *IEEE Trans. Haptics*, vol. 4, no. 4, pp. 321–326, Jul./Aug. 2011.
- [14] J. M. Hollerbach and C. W. Wampler, "The calibration index and taxonomy for robot kinematic calibration methods," *Int. J. Robot. Res.*, vol. 15, no. 6, pp. 573–591, Dec. 1996.
- [15] B. Knoerlein and M. Harders, "Comparison of tracker-based to tracker-less haptic device calibration," in *Proc. World Haptics Conf.*, Istanbul, Turkey, 2011, pp. 119–124.
- [16] G. Bianchi, B. Knoerlein, G. Szekely, and M. Harders, "High precision augmented reality haptics," in *Proc. Eur. Haptics Conf.*, Jun. 2006, pp. 169–178.
- [17] M. Huber, M. Schlegel, and G. Klinker, "Temporal calibration in multisensor tracking setups," in *Proc. Workshop der GI-Fachgruppe Virtuelle Erweiterte Realität*, Sep. 2012, pp. 201–212.
- [18] R. Y. Tsai and R. K. Lenz, "A new technique for fully autonomous and efficient 3D robotics hand/eye calibration," *IEEE Trans. Robot. Autom.*, vol. 5, no. 3, pp. 345–358, Jun. 1989.
- [19] M. Tuceryan, D. S. Greer, R. T. Whitaker, D. E. Breen, C. Crampton, E. Rose, and K. H. Ahlers, "Calibration requirements and procedures for a monitor-based augmented reality system," *IEEE Trans. Vis. Comput. Graph.*, vol. 1, no. 3, pp. 255–273, Sep. 1995.
- [20] B. K. P. Horn, "Closed-form solution of absolute orientation using unit quaternions," *J. Opt. Soc. Am.*, vol. 4, no. 4, pp. 629–642, 1987.
- [21] J. A. Hartigan and M. A. Wong, "Algorithm AS 136: A K-means clustering algorithm," *J. Roy. Statistical Soc. Ser.*, vol. 28, no. 1, pp. 100–108, 1979.
- [22] S. Maneewongvatana and D. M. Mount, "It's okay to be skinny, if your friends are fat," *Center Geometric Comput. 4th Annu. Workshop Comput. Geometry*, vol. 2, pp. 1–8, 1999.
- [23] U. Eck, F. Pankratz, C. Sandor, G. Klinker, and H. Laga, "[DEMO] comprehensive workspace calibration for visuo-haptic augmented reality," in *Proc. 17th IEEE Int. Symp. Mixed Augmented Reality*, 2014, pp. 333–334.
- [24] P. T. Boggs and J. E. Rogers, "Orthogonal distance regression," *Contemporary Math.*, vol. 112, pp. 183–194, 1990.
- [25] M. Huber, D. Pustka, P. Keitler, F. Echter, and G. Klinker, "A system architecture for ubiquitous tracking environments," in *Proc. 6th IEEE ACM Int. Symp. Mixed Augmented Reality*, Nara, Japan, 2007, pp. 1–4.



**Ulrich Eck** received the BSc (Honors) degree in computer science in 2011 from the University of South Australia after a professional career as an IT consultant, a software developer, and a business manager. He is currently working toward the PhD degree in the field of visuo-haptic augmented reality at the University of South Australia. His research interests include: medical applications, realtime systems, human-computer interaction, and software engineering. He is a student member of the IEEE.



**Gudrun Klinker** studied computer science (informatics) at the Universities of Erlangen and Hamburg. She received the PhD degree from Carnegie Mellon University in 1988. Since 1989, she has worked on topics of scientific visualization and augmented reality at Digital Equipment Corporation in Cambridge, MA, at ECRC in Munich and at the Fraunhofer Institute for Computer Graphics (FhG-IGD). Since 2000, she has been a professor for augmented reality at the Technische Universität München (TUM). She is a member of the IEEE.



**Frieder Pankratz** received the MSc degree in computer science in 2009 from the Technical University of Munich (TUM). He is currently working toward the PhD degree in the field of tracking management for augmented reality applications at the Technical University of Munich. His research interests include: augmented reality games, 3D user interfaces, and tracking.



**Hamid Laga** received the MSc and PhD degrees in computer science from Tokyo Institute of Technology in 2003 and 2006, respectively. He is currently a senior research fellow at the Phenomics and Bioinformatics Research Centre (PBRC), University of South Australia (UniSA). Prior to joining UniSA, he was an associate professor at the Institut Telecom, Telecom Lille1 in France (2010-2012), an assistant professor at Tokyo Institute of Technology (2006-2010), and a post-doctoral fellow at Nara Institute of Science



**Christian Sandor** received the doctorate in computer science from the Munich University of Technology, Germany. He is an associate professor at Nara Institute of Science and Technology, Japan. He was previously at Canon's Leading-Edge Technology Research Headquarters in Tokyo, Japan. Since the year 2000, his foremost research interest is augmented reality, as he believes that it will have a profound impact on the future of mankind. His further research interests include: human-computer interaction, visualization, computer graphics, computer vision, haptics, and software engineering. He is a member of the IEEE.

Technology in Japan (2006). His research interests span various fields of computer vision, computer graphics, and pattern recognition, with a special focus on the 3D acquisition, modelling and analysis of static and deformable objects. His contributions in these fields received the Best Paper Award at the IEEE International Conference on Shape Modeling (2006), the Best Paper Award at NICOGRAPH Paper Context (2007), the International Paper Grand Prix (Best Paper Award) by the Japan Society of Art and Science (2008), and the APRS/IAPR Best Paper Prize at the IEEE International Conference on Digital Image Computing, Techniques and Applications (DICTA) 2012.

tion, computer graphics, computer vision, haptics, and software engineering. He is a member of the IEEE.

▷ **For more information on this or any other computing topic, please visit our Digital Library at [www.computer.org/publications/dlib](http://www.computer.org/publications/dlib).**