

HapticHead: A Spherical Vibrotactile Grid around the Head for 3D Guidance in Virtual and Augmented Reality

Oliver Beren Kaul
University of Hannover
Hannover, Germany
kaul@hci.uni-hannover.de

Michael Rohs
University of Hannover
Hannover, Germany
michael.rohs@hci.uni-hannover.de

ABSTRACT

Current virtual and augmented reality head-mounted displays usually include no or only a single vibration motor for haptic feedback and do not use it for guidance. We present HapticHead, a system utilizing multiple vibrotactile actuators distributed in three concentric ellipses around the head for intuitive haptic guidance through moving tactile cues. We conducted three experiments, which indicate that HapticHead vibrotactile feedback is both faster (2.6 s vs. 6.9 s) and more precise (96.4 % vs. 54.2 % success rate) than spatial audio (generic head-related transfer function) for finding visible virtual objects in 3D space around the user. The baseline of visual feedback is – as expected – more precise (99.7 % success rate) and faster (1.3 s) in comparison, but there are many applications in which visual feedback is not desirable or available due to lighting conditions, visual overload, or visual impairments. Mean final precision with HapticHead feedback on invisible targets is 2.3° compared to 0.8° with visual feedback. We successfully navigated blindfolded users to real household items at different heights using HapticHead vibrotactile feedback independently of a head-mounted display.

Author Keywords

Guidance; navigation; haptic feedback; vibrotactile; virtual reality; augmented reality; spatial interaction; 3D output

ACM Classification Keywords

H.5.2. Information interfaces and presentation: User Interfaces – haptic I/O, input devices and strategies

INTRODUCTION

Navigation and 3D guidance systems use a large variety of different technologies to stimulate the visual, auditory, or haptic channels. The visual channel is usually the channel of choice as it typically has a higher bandwidth than the other channels [29]. However, sometimes the visual channel is not the desired primary channel to be used for some

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

CHI 2017, May 06-11, 2017, Denver, CO, USA

© 2017 ACM. ISBN 978-1-4503-4655-9/17/05...\$15.00

DOI: <http://dx.doi.org/10.1145/3025453.3025684>

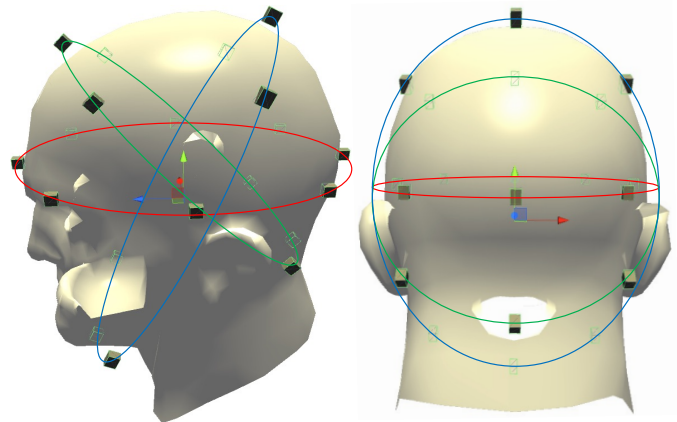


Figure 1. Placement of actuators in HapticHead. Note the three concentric ellipses around the user’s head and no actuators close to the ear openings. The red ellipse contains 8 equidistant actuators, the green and blue ellipses each contain 6 actuators.

kinds of feedback or in special situations such as when driving a car [6]. The visual channel might be overloaded and important feedback can be overlooked or lighting conditions may prevent the user from seeing the feedback at all. Another reason to use the tactile instead of the visual or auditory feedback channels are faster initial reaction times, as shown in several studies such as [25].

To relieve the visual channel, we propose HapticHead, a high-resolution, omnidirectional vibrotactile display worn on the head that presents 3D directional and distance information through moving tactile cues and patterns. It consists of a grid of vibrotactile actuators arranged in three concentric ellipses around the head for uniform coverage, optimized for head shape and user comfort (Figure 1). The head is well suited for guidance applications and tactile feedback, as it is sensitive to mechanical stimuli [9,18] and provides a large spherical surface. This allows displaying precise 3D information and allows the user to *intuitively* turn the head in the direction of a stimulus. We left important parts of the face uncovered and did not place actuators too close (less than 4 cm) to the ear openings because noise through bone conduction increases dramatically in their proximity.

HapticHead may be combined with virtual reality (VR) or augmented reality (AR) head-mounted displays (HMDs) to increase the sense of immersion. The visual content of the

HMD is then synchronized with vibrotactile feedback of HapticHead and the user's position and orientation. Users may see and feel shockwaves, particles, virtual walls, or moving virtual objects. As an example, "virtual bee" may "fly around", causing airflow, and even "sting" the user. This may be particularly effective when combined with spatial audio [31].

Beyond an increased sense of immersion, involving additional modalities also promises benefits in terms of user performance. Consider a VR or AR game with an already overloaded HMD due to lots of relevant data. HapticHead can be used to show or amplify critical directional and distance information, such as the direction of enemy fire.

There are several application possibilities of HapticHead in conjunction with HMDs. We here just hint at a few. Possible applications for 3D vibrotactile feedback around the head include: enemy positions and distances around fighter jet pilots; plane positions and distances around air traffic controllers; drone positions around drone pilots; fish and obstacle positions around scuba divers; star positions above stargazers; turn-by-turn navigation information for bikers; slope conditions ahead or approaching traffic behind skiers (collision feed-forward); object positions around crane operators; 3D guidance for visually impaired people; and 360° videos with haptic feedback.

Instead of exploring any of these scenarios in greater detail, the focus of this paper is on characterizing the performance of haptic guidance in 3D using a vibrotactile grid on the head, which is fundamental for most of the mentioned scenarios. We present the HapticHead concept and prototype, a guidance algorithm, and three experiments that evaluate different aspects and use cases of our system. The first experiment is designed to compare vibrotactile feedback vs. visual and auditory feedback in a virtual-object selection task. We then modified our prototype to increase precision and conducted a second experiment, comparing the new prototype vs. the old one and also the achievable precision of vibrotactile feedback vs. the baseline of visual feedback. The third experiment, which is independent of the others, shows that HapticHead feedback can be used in a real-world scenario for finding physical objects in a lab.

RELATED WORK

In the area of vibrotactile feedback for guidance and navigation or spatial awareness there are several works on haptic shoes, belts, bracelets, hand- and head-worn devices.

Paneels et al. [21] investigate tactile patterns on a tactile bracelet for indicating directions and find that static patterns are not well recognized due to the actuators being too close and being recognized as one impulse instead of multiple (phantom sensation) while dynamic patterns are recognized with higher accuracy. Weber et al. [30] investigate auditory vs. tactile stimuli through a wristband for guiding the arm and show that both have a similar performance for a translational task.

ActiveBelt [28] is the first vibrotactile belt for directional navigation. Van Erp et al. [7] use a vibrotactile belt for waypoint navigation. Vibrotactile belts have also been used to increase the situational awareness of gamers [22] and for guiding visually impaired people [4]. Meier et al. [16] investigate pedestrian navigation and compare three different vibrotactile devices (wristband, shoes, and belt) feedback and conclude that vibrotactile feedback alone might not be sufficient in complex geographic situations.

Haptic Radar [3] is a ring around the head, consisting of multiple infrared sensors and vibrotactile actuators in order to give users a "spider sense" of approaching objects. A similar concept is Proximity Hat [1], using pressure instead of vibrotactile actuators, which stimulates other receptors. Kerdegari et al. [13] developed a firefighter helmet with seven vibrotactile actuators on the forehead. Their experiment shows lower route deviation in a navigation task for vibrotactile compared to auditory feedback.

Israr et al. [10] present "Tactile Brush", an interpolation concept for multiple tactile actuators arranged in a grid in order to purposefully generate a moving tactile phantom sensation, which simulates the feeling of a continuous motion with a single localization point even though multiple actuators are active at a time. Schneider et al. [24] compared different interpolation strategies for the phantom sensation on 2D grids and found that logarithmic interpolation was rated higher than linear interpolation for a moving stimulus along a straight line. Further work on the vibrotactile phantom sensation [14] shows that the spacing of actuators should be 2.5 cm or less on the forehead for the phantom sensation to occur for most users. HapticHead uses a spacing of at least 4 cm (second prototype), so we do not expect the phantom sensation to arise. Myles et al. [19] investigate the vibrotactile sensitivity of different head regions and use a headband with 4 actuators to provide navigational cues to soldiers. They found that soldiers preferred a tactile to a visual or auditory display for directional cueing and that the forehead, frontal, parietal and temple regions were most sensitive to tactile stimuli.

Dobrzynski et al. [5] investigate information transfer capabilities of a ring of 12 vibrotactile actuators around the head with regard to the maximum number of active actuators and the maximum comfortable vibration intensity. They strongly suggest avoiding the use of multiple simultaneous actuators to show different directions.

All of the discussed related publications have in common that they do not *intuitively* guide in three dimensions as they only use a single ring of vibrotactile actuators or less and thus can only map signals on a 2D plane with distance-related vibrotactile patterns. Conceptually, Cassinelli et al. [3] discuss extensions of their ring prototype and propose to place modules anywhere on the body but as far as we know did not implement or test this.

INITIAL PROTOTYPE

Myles and Kalb [19] recommend actuators on the head to operate at frequencies between 32 and 150 Hz because of discomfort above that threshold. We decided to use actuators operating at 150 Hz at maximum because actuator size increases for equally strong impulses at lower frequencies. Our first prototype (Figure 2) consists of a bathing cap with 17 vibration motors (Parallax, 12 mm coin type, 3.3 V, 90 mA, 9000 rpm) attached on the inside (Figure 2, left). The non-stretchable chinstrap hosts another three vibration motors and can be adjusted to different head sizes using a Velcro fastener. The vibration motors are controlled by PWM signals of four Arduino Nanos on a switchboard, which are connected to a stationary PC through USB and are updated at 75 Hz.

On the software side, vibration motors are modeled at their corresponding position in a Unity [32] scene. This allows easy spatial activation of selected motors, depending on the task. The user’s head is tracked either by the internal sensors of an HMD to be used in conjunction with HapticHead or by an external tracking system such as OptiTrack.

EXPERIMENTS

We conducted three experiments in order to characterize user performance and to refine our 3D guidance concept in virtual and real environments. In all experiments targets may be located at any position around the head. This includes positions that are not in the visual field initially and positions above and below the user.

Experiment 1 is a follow-up of the experiment in [12], using HapticHead combined with an Oculus Rift DK2 to find visible virtual targets equally distributed on a sphere around the user. In the previous experiment we found indications that HapticHead feedback might be an interesting alternative to visual and auditory feedback but did not evaluate this in detail. The current experiment includes more participants and additionally records movement trajectories and success rates individually per target. Based on the results, the participants’ comments, and our observations, we refined the concept and prototype. Experiment 2 evaluates performance differences due to refinements of the prototype and the guidance algorithm. Furthermore, the achievable precision with both visual (attention funnels and one-pixel targets) and vibrotactile (invisible targets) feedback is investigated.

Experiment 3 is independent of the other experiments and aims to show the usefulness of our concept for finding tracked physical objects around blindfolded users, i.e., without visual feedback.

EXPERIMENT 1: VIRTUAL VISIBLE TARGETS

As a first step we evaluate the performance of HapticHead in guiding users who wear an Oculus Rift DK2 towards virtual 3D objects around the head (Figure 2, right). This allows us to then refine the concept and prototype based on the findings.




Figure 2. First HapticHead prototype (left) and Unity scene view with visible targets from the outside (right). The user’s camera is in the center.

We built a simple VR environment in Unity 5.3 that spawns 20 small ($r = 1$ m) equidistant spheres on the surface of a larger ($r = 5$ m) invisible sphere with the viewer at its center (see Figure 2). The spheres were distributed with pack.3.20 coordinates [27]. As the user rotates the head the location of the spheres stays fixed with respect to the environment. The target spheres do not coincide with the actuator positions.

There are three feedback conditions in Experiment 1: *visual*, *auditory*, and *vibrotactile* feedback. We included visual feedback as a baseline because AR and VR applications are usually designed around visual feedback. We also included auditory feedback, as auditory feedback is often used when applications aim not to overload the user’s visual sense.




Figure 3. Attention funnels with a tiny red crosshair in the view’s center. Visual feedback from the user’s perspective.

In the *visual feedback* condition guidance towards objects is achieved through the concept of *attention funnels* as in [2] which are a state-of-the-art 3D visual guidance concept utilizing “target goals” to guide a user. We implemented attention funnels using the same green target goals as in [2] and made sure that they also work with targets behind the user (Figure 3).

For the *auditory feedback* condition we used white noise in combination with Unity 5.3’s included spatial sound system (which uses a generic head-related transfer function, g-HRTF) with “spatial blend” set to 1 (full 3D) and Bose QC25 stereo noise cancellation headphones (NC off). We are aware that there are better audio technologies available which utilize personalized head-related transfer functions (p-HRTFs) but these require a complex per-user calibration. HapticHead does not require per-user calibration.

In the *vibrotactile feedback* condition, we activate the three actuators that are closest to the target with an interpolated intensity that represents closeness. The closest actuator is

running at the highest intensity. Actuator intensities are not static and are rather adjusted with head rotation. So as the user turns the head towards the target, the signal travels along the trajectory towards the front of the head.

We invited 13 participants (2 female, mean age 23.5, SD 3.2 years). Only 5 had previous experience with VR HMDs.

Participants had to focus a “start box” at a fixed position for half a second in order to start a trial. Feedback was turned on and the task was to find the target as quickly and accurately as possible. Once the participants had located and focused the suspected target, they pressed a hand-held button to end the trial. Upon pressing the button, the sphere visually highlighted in green or red, depending on whether it was the right one. Instead of dedicated training trials, we chose this form of active highlighting to measure any learning effects and whether participants would “calibrate themselves” towards this new form of haptic feedback.

Each participant performed 480 trials: 3 feedback conditions (visual, auditory, vibrotactile) \times 20 targets \times 8 repetitions per target. The three feedback conditions were presented in blocks. Their order was counterbalanced with a Latin square. The order of targets was random. As dependent variables we measured the head movement trajectory, task completion time, and error rate. The user’s focus point was tracked with each frame, at a rate of 75 Hz, which was also the update rate for all feedback conditions and data logging. Participants could pause between each trial and had a forced pause when the feedback condition changed. The experiment took around one hour per participant. As a reward each participant received a bar of chocolate.

Results of Experiment 1

While running the experiment we observed randomly appearing frame lags of 1 s + frame time due to a graphics driver problem. Because of this we had to exclude 248 trials (3.97 %). In order to reduce the influence of outliers on task completion time, we used the median of the 8 repetitions per target that each user performed.

We define *movement overhead* (m.o. in Figs.) as the ratio of the actual path length to the optimal path length, minus 1,

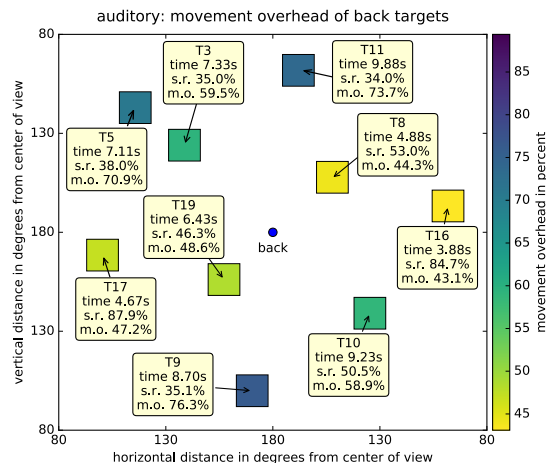
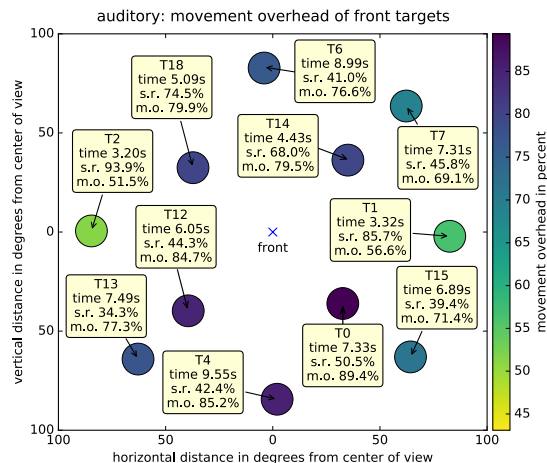


Figure 4. Auditory condition, front (left) and back (right) targets: mean movement overhead (m.o.), median trial completion times and success rates (s.r.)

i.e., the movement overhead is a percentage on how much longer the user’s trajectory is compared to the optimal path on the sphere towards the target. A value of 0 % means the user exactly follows the optimal path. A value of 100 % means the user’s trajectory is twice as long as the optimal path. Movement overhead thus is a measure of how directly the user is able to localize the target.

	Median trial time [s]	Mean trial time [s] (SD)	Mean movement overhead [%] (SD)	Success rate [%]
Visual (att. funnels)	1.22	1.28 (0.37)	8.8 (3.6)	99.66
Auditory (g-HRTF)	6.28	6.86 (3.55)	67.2 (14.4)	54.22
Vibrotactile (HapticHead)	2.41	2.61 (1.04)	35.1 (10.7)	96.36

Table 1. Task completion times and success rates for different feedback conditions.

Table 1 and Figure 5 show the measured dependent variables with merged data from all participants and all trials, not just successful ones.

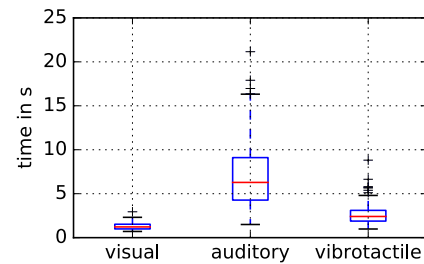


Figure 5. Boxplots of completion times for all conditions with merged data from all participants.

The auditory condition was the slowest at 6.86 s. At 2.61 s vibrotactile only took 38 % and at 1.28 s visual feedback only took 19 % of the time of auditory feedback. A two-way repeated-measures ANOVA shows statistically significant main effects of feedback condition ($F_{2,24} = 50.30, p < 0.001$) and target ($F_{19,228} = 12.25, p < 0.001$) on task completion time, and an interaction effect of feedback condition and target ($F_{38,456} = 5.06, p < 0.001$). A Friedman test reveals a significant difference in success rates between

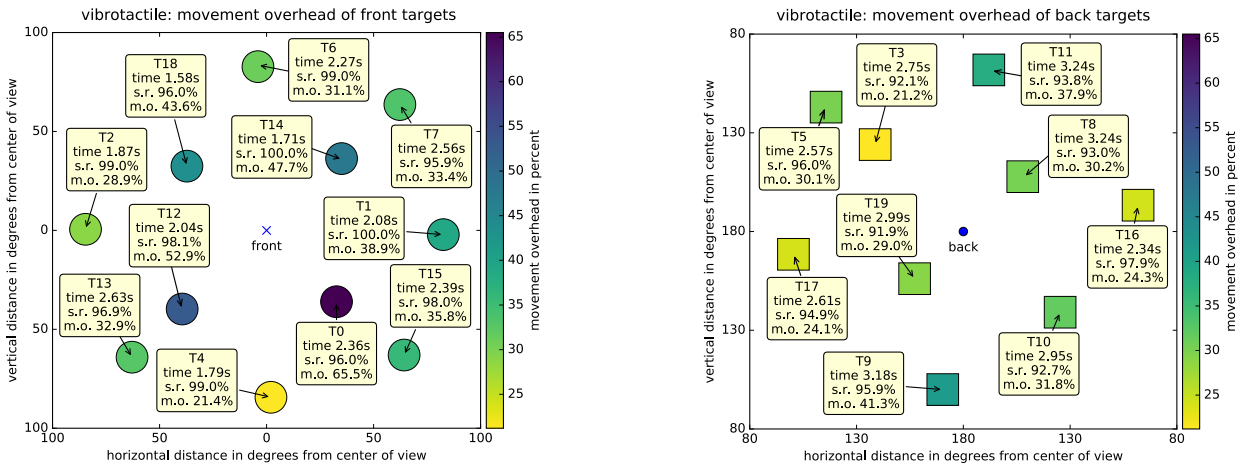


Figure 6. Vibrotactile condition, front (left) and back (right) targets: mean movement overhead (m.o.), median trial completion times and success rates (s.r.)

conditions ($\chi^2(2) = 25.04, p < 0.001$). Individual success rates (s.r.) are shown in Figures 5 and 6.

Results for movement overhead are very much as expected for the visual feedback condition (attention funnels), with nearly perfect and immediate localization of the target. Therefore we do not include a detailed report of these results. For comparison, the mean movement overhead for all targets with visual feedback was 8.8 % (SD 3.6 %).

In Figures 5 and 6, all targets are projected onto a vertical plane for the initial state of the user looking at the starting cube. The vertical plane's normal vector is horizontal and points in the user's frontal direction.

The auditory feedback condition had a mean movement overhead of 67.2 % (SD 14.4 %). Figure 4 shows targets in front and back of the user for the auditory condition. Please note that the color scales are different between modalities. For the auditory (g-HRTF + white noise) condition, guidance towards targets near the horizontal plane through the ears works much better than towards targets off the horizontal plane. The dependent variables show a clear trend for targets near the horizontal plane being faster to reach, with a higher success rate and a lower movement overhead. Targets directly above and below the user were particularly hard to find and took users a long time to identify.

The vibrotactile feedback condition had a mean movement overhead of 35.1 % (SD 10.7 %). As Figure 6 shows, compared to auditory feedback, targets off the horizontal plane worked much better with vibrotactile feedback. The performance was higher for all the measured variables.

However, targets such as T0, T12, and to a lesser extent T9, T14, and T18 had unexpectedly higher-than-average movement overheads. We presume that this is due to the chin belt of the first prototype being too inflexible and distributing vibrotactile signals from one actuator along the whole chin (T0, T12), and also due to our guidance algorithm interpolating between the three actuators closest to the target which turned out not to be the best solution for

non-uniform actuator distributions. We describe solutions to these issues in the next section.

Comparing the mean movement overheads for front targets on a diagonal guidance path (T0, T7, T12, T13, T14, T15, T18; mean movement overhead: 44.5%) to those on vertical or horizontal guiding paths (all other targets; mean movement overhead: 30.1%) suggests that users had more problems locating targets on the diagonals.

Vibrotactile targets on the front had an average completion time of 2.12 s, targets on the back of 2.87 s. This is expected because users first need to turn around to reach targets behind them. Visual targets in comparison had an average time of 1.01 s on the front and 1.54 s on the back. The average success rate of front targets in the vibrotactile condition was 98.0 %, whereas it was 94.2 % for back targets.

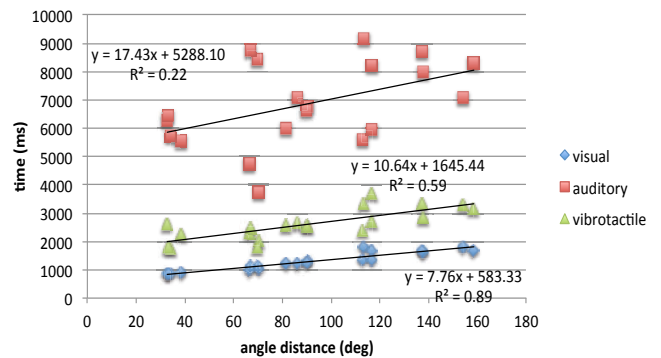


Figure 7. Selection time by angular distance between the start box and the target center. Visual ($R^2 = 0.89$) and vibration ($R^2 = 0.59$) conditions show a linear relationship.

Figure 7 shows the selection time by angular distance between the starting orientation of the user and the orientation of each target. For the visual ($R^2 = 0.89, p < 0.001$) and vibrotactile ($R^2 = 0.59, p < 0.001$) conditions there is a good linear fit between angular distance and selection time. For the auditory condition there is a much weaker relationship ($R^2 = 0.22, p < 0.001$). For each target, the visual and vibrotactile conditions outperform the auditory condition.

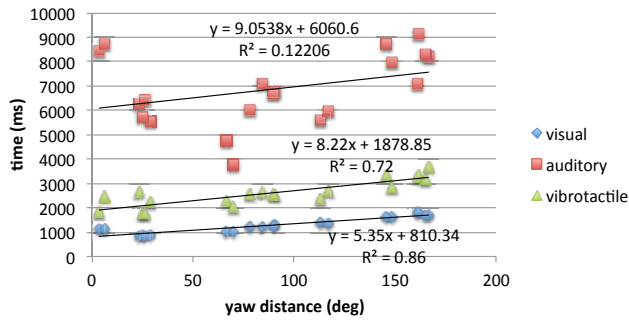


Figure 8. Selection time by yaw (horizontal heading) distance between the start box and the target center. Visual ($R^2 = 0.86$) and vibrotactile ($R^2 = 0.72$) conditions show a linear relationship. The auditory condition is slow with targets in front of and behind the user.

Figure 8 shows the selection time by yaw (horizontal heading) distance between the starting orientation and each target. Again, there is a good linear fit for the visual ($R^2 = 0.86$, $p < 0.001$) and vibrotactile ($R^2 = 0.72$, $p < 0.001$), but not for the auditory condition ($R^2 = 0.12$, $p = 0.025$).

Our observations show that in the auditory condition participants had trouble locating targets below or above ear level. The further targets are off the horizontal plane at ear height, the longer the selection time. This effect is symmetric above and below the plane and roughly has a parabola shape. The visual and vibrotactile conditions do not exhibit this effect.

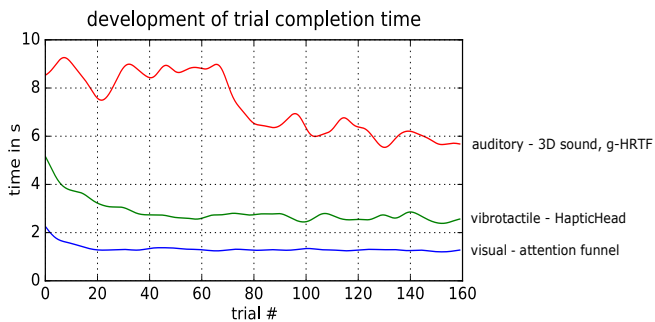


Figure 9. Average learning effect – time. Merged data, all participants. Curves: Gaussian weighted moving average (width=3, blue=visual, green=vibrotactile, red=auditory).

Figure 9 shows the development of completion time over all trials. Note the steep learning curve for vibrotactile feedback that flattens around trial 40 and the auditory learning curve, which initially stays rather constant until around trial 65 when participants started being quite a bit faster.

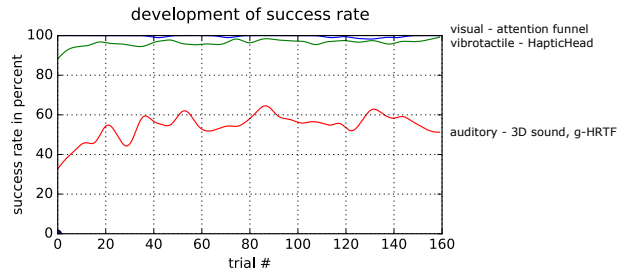


Figure 11. Average learning effect - success rate. Merged data, all participants. Curves: Gaussian weighted moving average (width=3, blue=visual, green=vibrotactile, red=auditory).

Figure 11 shows the success rates over time. In the auditory condition, participants had a steep learning curve, which flattens after trial 40 but still shows a high variance compared to the other conditions. Despite the fact that participants learned to be a lot faster after trial 65 in the auditory condition, the success rate did not drop but even increased a bit. In the vibrotactile condition, participants needed less than 15 trials to accommodate themselves with this new form of feedback. After the first few trials, the success rate curve for vibrotactile feedback flattens and stays close to 97% without much variance or measurable fatigue effects.

Qualitative Results

Qualitative results were measured through a post-questionnaire with 5-point Likert scales (Figure 10). Participants agreed that the vibrotactile feedback was helpful for finding virtual objects while they disagreed that the auditory feedback was helpful. They agreed that the feedback position around the head was appropriate and that the vibrotactile feedback was comfortable. Participants had mixed opinions about the sound level of the vibrotactile feedback. They agreed that the vibrotactile feedback was unambiguous when looking at the correct target and most participants could imagine using such feedback regularly.

Discussion of Experiment 1

Our experiment clearly shows that HapticHead vibrotactile feedback guides users towards visible virtual targets around

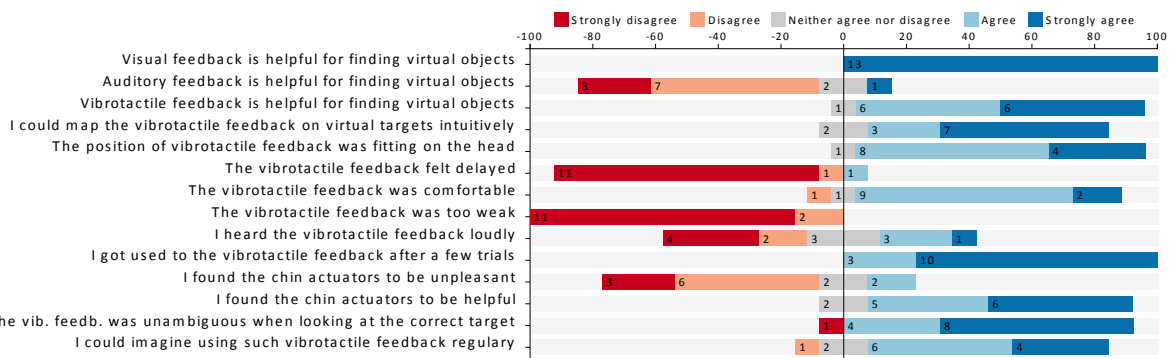


Figure 10. Qualitative results of Experiment 1. Diverging stacked bar chart: scales in percent, and absolute values on bars.

them substantially faster than spatial auditory feedback (directional hearing with a g-HRTF). For vibrotactile feedback there is a linear relationship between angular distance and selection time. In addition to higher selection times, weak points of g-HRTF auditory feedback include confusion between targets directly in front of and behind the user as well as issues in locating targets above or below the horizontal ear plane, which supports earlier work [17]. This is expected because g-HRTF auditory feedback is known to cause localization difficulties and can be improved by using p-HRTFs [8] or letting users do training on g-HRTF localization with visual feedback on the right target as in the work by Klein et al. [15]. We also saw a large improvement in performance for the auditory condition as users adapted to the g-HRTF due to getting visual feedback on the correct target after each trial. Fortunately, vibrotactile feedback does not share these issues and more closely resembles the performance characteristics of visual feedback. Vibrotactile feedback thus suggests itself as an alternative to visual feedback. However, we could not measure the achievable precision due to the arbitrary target size and we saw opportunities to refine our prototype and guidance algorithm.

REFINEMENT OF CONCEPT AND PROTOTYPE

Based on the experiences from Experiment 1 we built a second prototype in order to improve precision for guidance applications and user comfort. As discussed in the previous section, the non-stretchable chin belt of our first prototype seemed to have a substantial negative impact on guidance towards targets on the diagonals below the user (T0, T12). For the second prototype we replaced the chin belt with a stretchable one. We also increased the number of actuators on the forehead and on the chin belt from three to five in order to form a “ring” of actuators around the face. We made sure to avoid the ear openings to minimize noise through bone conduction.

For the second prototype we placed the actuators on the outside of the bathing cap, this time due to feedback from experiment participants who commented on the vibration



Figure 12. Second, refined HapticHead prototype, side and front view. Notice actuators located on the outside, the flexible chinstrap, and five instead of three actuators on each, the forehead and chinstrap. Positions of the 10 forehead and chin actuators forming a “ring” around the face are marked in red.


motors leaving tiny marks on the forehead. For the motors on the chinstrap however we could not place them on the outside because the vibrotactile impulse would have been attenuated too much otherwise.

In order to make our prototype untethered we exchanged the four Arduinos with a single Raspberry Pi 2 on a custom amplification board with a Wi-Fi dongle and a standard 5 V USB battery pack. For the vibration motors we used a total of 24 Precision Microdrives 310-117 - Pico Vibe [23] (10 mm diameter x 3 mm height, 150 Hz frequency at 3.3 V, low starting voltage of 0.9 V, controlled at 500 Hz software PWM by the Raspberry Pi 2, updated at 75 Hz).

We also refined the guidance algorithm, whose aim is to adapt the intensity of the actuators so as to best guide the user towards a target. The previous version of the algorithm just interpolated between the three actuators closest to the target. This can cause unintuitive behavior especially when the target is right in front of the user’s face.

For the refined guidance algorithm we defined a virtual point zero (“VPZ”) exactly between the eyes of the user (marked in green in Figure 12 and Figure 13). We then tessellated the actuator space in that we placed triangles between each triple of adjacent actuators (including the VPZ) without overlaps as shown in Figure 13.

A ray between the center of the head and the virtual target around the head intersects exactly one of the triangles t of the tessellation in a hit point h . Triangle t is defined by its adjacent actuators (v_0, v_1, v_2). Let point e_i be the intersection of a line through v_i and h and a line through $v_{(i+1) \bmod 3}$ and $v_{(i+2) \bmod 3}$, the other two actuator positions (\bmod is the modulo operation).



The intensities (0 to 1) of the three actuators are calculated as:

$$intensity(v_i) = 1 - \frac{|h - v_i|}{|e_i - v_i|}$$

Figure 14. Intensity calculation visualization.




Figure 13. Side and front view of modeled actuator positions. Does not fit perfectly due to arbitrary size and asymmetries of the Styrofoam head. The refined guidance algorithm uses triangles between actuators, including a virtual point zero (in green) between the eyes.

There is a special case: If the VPZ is part of the intersected triangle, the intensity of the remaining two actuators is amplified in order to give the user a sense of direction on the “ring around his face”. The user is then drawn a bit more in the indicated direction.

This way of computing the intensities is similar to the 2D linear approach in the work by Schneider et al. on “Tactile Animation” [24] whose experiment however indicated that users rated a 2D straight-line motion best using a logarithmic approach over a linear one. On the other hand, Seo et al. [26] found that location accuracy in a 1D case between two actuators was higher with a linear approach than with a logarithmic one. Since we focus on localization precision, we chose a linear interpolation approach. Our 3D-sphere linear interpolation approach achieves several objectives: The experienced overall tactile intensity is only weakly dependent on the position of the hit point in the triangle, just like in [24]. The algorithm worked well in pilot testing despite non-uniform placements of the actuators. Whether the same algorithm with a logarithmic interpolation would outperform our chosen linear interpolation in a 3D scenario remains an open topic for future work.

We also included a vibrotactile pattern to indicate angular closeness of the VPZ to the target. In a variant (Experiment 3) we mapped the pattern to the depth-axis instead. Actuators have a duty cycle of 95% for a 1 Hz pattern if the user is above a certain distance threshold (either angular or depth). A 95% duty cycle instead of 100% was chosen to periodically remind the user that the pattern is still there while the distance is still way off. Once below the threshold, the duty cycle linearly decreases while the frequency increases until the actuators have a 70% duty cycle for a 50 Hz pattern at zero distance. This is experienced as permanently on with an intensity of 70%, as the motor’s stop time (from full speed to stop) is greater than the signal off-time. Thresholds should be chosen depending on the task. Choosing a larger threshold results in less time needed to complete a task but also less accuracy. In pilot testing these values and the resulting vibrotactile pattern were found to be effective but remain to be optimized in future work.

EXPERIMENT 2: VIRTUAL INVISIBLE TARGETS

We were interested in how *precise* users could be with the visual and vibrotactile feedback mechanisms and how our second prototype and new guidance algorithm would perform compared to the first prototype. As in Experiment 1, users wore an Oculus Rift DK2 in all conditions and data logging and refresh rate of actuators was the same.

Experiment 2 uses a within-subject design with feedback type as the independent variable. There are three levels for feedback condition: *vibrotactile-visible-targets*, *visual-1-pixel*, and *vibrotactile-invisible* targets. The *vibrotactile-visible-targets* condition is designed to be compared to the vibrotactile condition from Experiment 1 and not to the other conditions in this experiment.

Vibrotactile-visible-targets condition: For a comparison of the first prototype and old guidance algorithm we ran the *vibrotactile* feedback condition from Experiment 1 again with the new prototype and guidance algorithm (same target size, 8 repetitions \times 20 targets per user). The guidance algorithm used the pattern as explained above with an angular threshold of 40°, which was determined in a pre-experiment and remains to be optimized in future work.

Visual-1-pixel condition: Since final precision cannot be directly concluded from Experiment 1 due to the arbitrary target size, we decided to use visual feedback as a precision baseline condition. Attention funnels with a target crosshair as in Experiment 1 (shown in Figure 3) were used, but here with tiny white 1-pixel targets instead of the large ones (4 repetitions \times 20 targets per user).

Vibrotactile-invisible-targets condition: We used invisible targets with the refined prototype and the new guidance algorithm as our vibrotactile precision condition (8 repetitions \times 20 targets per user). The guidance algorithm used the vibrotactile pattern described above with an angular threshold of 16°, which was also determined in a pre-experiment. The angular threshold chosen for this condition is smaller than in the *vibrotactile-visible-targets* condition because this condition does not have visual target markers and we wanted to focus on maximum possible precision within an acceptable timeframe.

The trials were executed as in Experiment 1. The correct target was indicated by a green sphere ($r = 1.0$ m for the *vibrotactile-visible-targets* condition and $r = 0.2$ m for the others) at the correct position after each trial.

We invited 13 participants (3 female, mean age 22.8, SD 2.6 years). The set of participants for Experiment 2 was fully disjoint from the set of participants in Experiment 1, so no participant had prior experience with HapticHead.

Results of Experiment 2

The *vibrotactile-visible-targets* condition had a mean trial completion time of 4.30 s, which is 65% higher than in the previous study. An increase was expected because participants had to wait for the vibrotactile pattern in order to confirm a target. However, the success rate increased only marginally to 96.6%. We suppose that this is because we had one participant (ID 11) who focused on being fast instead of precise. With a mean trial completion time of 2.67 s, this participant was 61% faster than the average but only reached a success rate of 86.8%. The other participants had a mean success rate of 97.5%.

For the other two conditions (*visual-1-pixel* vs. *vibrotactile-invisible-targets*) a two-way repeated-measures ANOVA shows significant main effects for feedback condition ($F_{1,12} = 24.88$, $p < 0.001$) and target ($F_{19,228} = 8.70$, $p < 0.001$) on trial time, and a significant interaction of feedback condition and target ($F_{19,228} = 2.28$, $p < 0.01$) for trial time. Furthermore, feedback condition ($F_{1,12} = 14.69$, $p < 0.01$) and target ($F_{19,228} = 2.08$, $p < 0.01$) have statistically

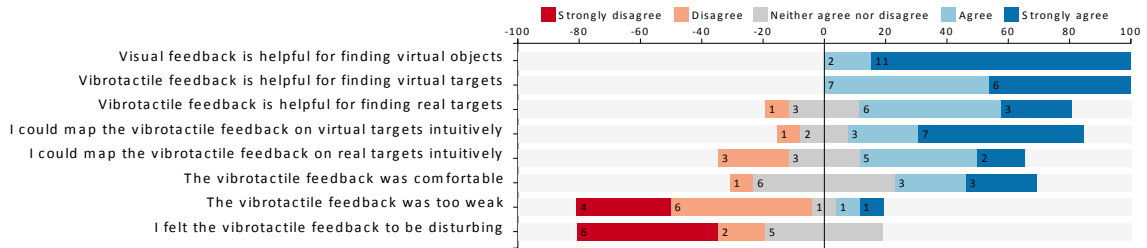


Figure 15. Qualitative results of Experiments 2 and 3. Diverging stacked bar chart: scales in percent, and absolute values on bars.

significant main effects on precision. There is also a significant interaction effect ($F_{19,228} = 2.25, p < 0.01$) for precision.

The main purpose of this study was to evaluate precision with the refined prototype and algorithm. In the *vibrotactile-invisible-targets* condition, participants reached a mean final deviation from the target of 2.33° (SD 1.80° , 95% CI [$0.59^\circ, 7.13^\circ$]) with a mean trial completion time of 8.92 s. The mean final deviation in the *visual-1-pixel* condition was 0.80° (SD 0.43° , 95% CI [$0.24^\circ, 1.80^\circ$]) with a mean trial completion time of 3.41 s. In comparison, the vibrotactile condition is less precise but still very close to the target.

Qualitative Results

As shown in Figure 15, participants agreed that the HapticHead vibrotactile feedback was helpful for finding virtual targets and most of the participants could intuitively map the feedback to the targets. Participants weakly agreed that the vibrotactile feedback was comfortable and disagreed about it being disturbing. They also mostly disagreed that they felt the vibrotactile feedback to be too weak.

Discussion of Experiment 2

While the time to find targets increased by more than what we expected in direct comparison to the old prototype and guidance algorithm, we also saw a small increase in success rates and measured only a small error of 2.3° towards targets in the *vibrotactile-invisible-targets* condition. With such a small error it should be easy to find targets in a real-world scenario, which leads us to our final experiment.

EXPERIMENT 3: REAL TARGETS

Independently of Experiments 1 and 2, we were interested if the HapticHead concept can also be used as sole feedback variant to find *real* targets around the user such as keys lost on top of the fridge or guiding visually impaired people into particular directions. *Research question 1*: Can blindfolded users find real objects around them with HapticHead vibrotactile feedback?

We used the OptiTrack tracking system for positions of real targets and for the position and orientation of the user’s hand. This time the user’s primary hand was the “center point of attention” because in pilot testing it proved to be more intuitive than the center of the head when trying to grab something. This means that the direction of the target object relative to the hand was displayed on the head. *Research question 2*: Is this indirection still intuitive for users? The effective position of the hand tracker was manually calibrated for each user to be on the palm of their hand. The guidance algorithm used a vibrotactile pattern as described above with a depth threshold of 1.5 m. This is different from Experiment 2 where an angular threshold was used to trigger the final phase and the vibrotactile pattern. Here, the pattern is used to show target depth instead because when trying to find real targets, it is important to also give users an impression of the remaining distance to the targets so they do not accidentally grab a knife with too much speed. The threshold distance was determined in a pre-experiment.

We conducted this experiment right after Experiment 2 with the same participants, which allowed us to skip training trials. In the lab ten small items were either hanging from the ceiling on small threads at different heights or were placed on a table or on the ground. The items on the table were three books among 10 books at 0.9 m height, so this was a “choose one from many” search task. Of the other items, two were placed on the floor and another two 1.7 and 1.8 m from the ground. The remaining three items were placed at comfortable heights between 0.8 and 1.4 m. This configuration of items was same for all participants. The items were equipped with OptiTrack markers (Figure 16).

The blindfolded participants started a trial in the center of the room within a small marked square, facing in the same direction every time. In each trial, their task was to find one randomly chosen item (no repetitions) at a distance of 1.2 to 2.5 m. In each trial the user was guided to one of the items using HapticHead vibrotactile feedback and the trial was



Figure 16. 10 items for Experiment 3. From left to right (height): Three books (0.9 m), a pen (1.8 m), a Lego piece (1.1 m), a ball (1.7 m, 6 cm diameter), a screwdriver (0.8 m), a remote control (1.4 m) and two balls (3 cm and 12 cm diameter, on the ground).

manually stopped when the participant was sure to have found the right target and said “stop”. The experimenter took note whether the target was the right one. We measured task completion time and whether participants found the right target. In addition, questionnaires gathered qualitative feedback on intuitiveness and usefulness.

Results of Experiment 3

Six of the seven non-book targets were found by all participants. Two participants (15.4 %) failed to find the last target, which was a small pen at a height of 1.8 m. They were both rather short people and remarked that they did not expect something that high up.

The correct book targets were found successfully in 52.6 % of all cases. However, the tracking method (OptiTrack with marker attached to the user’s hand, effective tracking position calibrated below the user’s palm) introduced tracking errors for the book targets due to the way some participants turned their hand while searching for these targets. Because we logged position and orientation of the hand marker we were able to manually classify failed trials where the hand marker actually pointed at the right target as successful. Excluding errors introduced by the tracking method, two out of three books were found successfully by all participants. The thinnest book was missed 3 out of 13 times. An overall success rate of 96 % was achieved over all targets. On average it took 42.0 s (SD 45.0 s) to find the correct target. Participants were rather cautious and moved slowly because they were not used to being blindfolded and did not want to run into things.

Qualitative Results

As shown in Figure 15, participants found the vibrotactile feedback helpful for finding real targets around them and could intuitively map vibrotactile signals to targets, thus research question 2 can be answered positively. Both of these measures were less agreed on than for virtual targets though. We believe this is because participants were used to the angular vibrotactile pattern (Experiment 3 used a depth pattern) and the algorithm reacting to position and orientation changes of their head instead of their hand. Because participants did only 10 trials each with these changes (no training trials) they had little time to get used to them.

Discussion of Experiment 3

Experiment 3 shows a high potential of the HapticHead concept to be used in a real world application in conjunction with a suitable tracking system for finding items or simply to orient a user in the right direction.

LIMITATIONS

Giving feedback on the head implies some anatomy limitations such as the need to keep a user’s face free of actuators and the thickness of user’s hair, which can weaken the stimulus received [20]. In our experiments, we had two participants with thick hair who indicated that they did not receive sufficiently strong feedback on the top of their heads. These participants needed more time to find the correct targets but had a similar success rate as the others.

We attribute this to the frontal vibrotactile actuators on the user’s forehead, which are unimpeded by hair.

Using Unity 5.3’s included audio system for comparison is a limitation as this system only uses g-HRTFs and is thus not a state-of-the-art audio system. However, p-HRTF systems require a complex per-user calibration, which HapticHead does not. Still, using a p-HRTF system would likely improve auditory results substantially and a comparison between HapticHead feedback and auditory p-HRTF feedback remains an interesting topic for future work.

Vibration causes noise through bone conduction. When asked whether they heard the vibrotactile feedback too loudly, participants had mixed opinions. However, this might be a more severe issue for visually impaired users as they rely on their sense of hearing more strongly than people with normal vision. This can be solved with actuators running at a lower frequency as suggested in [19] or an entirely different kind of actuator such as electro-tactile electrodes to generate a local tickling sensation as in [11]. Long term effects of this kind of tactile stimulation on the head are unknown but we do not expect strong effects as participants of our studies rated the device rather comfortable and our tactile stimulation is comparable to (strong) auditory bone conduction stimulation (20–150 Hz) on mostly soft tissue which dampens the effect. Bone conduction speakers are widely available and considered safe. However, a long-term study on tactile stimulation on the head remains an open topic for future work.

We did not investigate the influence of a potential phantom sensation, which some users might experience. If users feel one instead of two or three stimuli when two or three actuators are active at a time this could have an impact on performance and remains to be investigated in future work.

CONCLUSION

The results of the experiments show that a spherical grid of vibrotactile actuators around the head together with our guidance algorithm can effectively guide users towards virtual and real targets in 3D. Vibrotactile feedback turned out to be superior to g-HRTF auditory feedback and almost on par with visual feedback. The main contributions are the HapticHead concept itself, the guidance algorithm, and the evaluation of speed and precision of a prototype implementing the concept.

The concept itself can be used in a variety of other use cases, as hinted at in the introduction. We already implemented a few VR immersion example applications with HapticHead feedback and look forward to evaluate these in a qualitative experiment. In future work, we will also investigate a possible use of HapticHead for guidance and navigation to real objects for the blind and elderly people.

REFERENCES

- [1] Berning, M., Braun, F., Riedel, T., and Beigl, M. ProximityHat. *Proceedings of the 2015 ACM International Symposium on Wearable Computers* -

- ISWC '15*, ACM Press (2015), 31–38.
- [2] Biocca, F., Tang, A., Owen, C., and Xiao, F. Attention funnel: omnidirectional 3D cursor for mobile augmented reality platforms. *Proceedings of the SIGCHI conference on Human Factors in computing systems 6*, (2006), 1115–1122.
- [3] Cassinelli, A., Reynolds, C., and Ishikawa, M. Augmenting spatial awareness with haptic radar. *Proceedings - International Symposium on Wearable Computers, ISWC*, (2007), 61–64.
- [4] Cosgun, A., Sisbot, E.A., and Christensen, H.I. Evaluation of rotational and directional vibration patterns on a tactile belt for guiding visually impaired people. *IEEE Haptics Symposium, HAPTICS*, IEEE Computer Society (2014), 367–370.
- [5] Dobrzynski, M.K., Mejri, S., Wischmann, S., and Floreano, D. Quantifying Information Transfer Through a Head-Attached Vibrotactile Display: Principles for Design and Control. *IEEE Transactions on Biomedical Engineering* 59, 7 (2012), 2011–2018.
- [6] Van Erp, J.B.F. and Van Veen, H.A.H.C. Vibrotactile in-vehicle navigation system. *Transportation Research Part F: Traffic Psychology and Behaviour* 7, 4 (2004), 247–256.
- [7] Erp, J.B.F. Van, Veen, H. a. H.C. Van, Jansen, C., and Dobbins, T. Waypoint navigation with a vibrotactile waist belt. *ACM Transactions on Applied Perception* 2, 2 (2005), 106–117.
- [8] Geronazzo, M., Bedin, A., Brayda, L., Campus, C., and Avanzini, F. Interactive spatial sonification for non-visual exploration of virtual maps. *International Journal of Human-Computer Studies* 85, (2016), 4–15.
- [9] Gilliland, K. and Schlegel, R.E. Tactile Stimulation of the Human Head for Information Display. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 36, 4 (1994), 700–717.
- [10] Israr, A. and Poupyrev, I. Tactile brush. *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*, ACM Press (2011), 2019.
- [11] Kajimoto, H., Kanno, Y., and Tachi, S. Forehead electro-tactile display for vision substitution. *Proc EuroHaptics*, (2006), 11.
- [12] Kaul, O.B. and Rohs, M. HapticHead: 3D Guidance and Target Acquisition through a Vibrotactile Grid. *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems - CHI EA '16*, ACM Press (2016), 2533–2539.
- [13] Kerdegari, H., Kim, Y., and Prescott, T.J. Head-Mounted Sensory Augmentation Device: Comparing Haptic and Audio Modality. In Springer International Publishing, 2016, 107–118.
- [14] Kerdegari, H., Kim, Y., Stafford, T., and Prescott, T.J. Centralizing bias and the vibrotactile funneling illusion on the forehead. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Springer Verlag (2014), 55–62.
- [15] Klein, F. and Werner, S. Auditory Adaptation to Non-Individual HRTF Cues in Binaural Audio Reproduction. *Journal of the Audio Engineering Society* 64, 1/2 (2016), 45–54.
- [16] Meier, A., Matthies, D.J.C., Urban, B., and Wettach, R. Exploring vibrotactile feedback on the body and foot for the purpose of pedestrian navigation. *Proceedings of the 2nd international Workshop on Sensor-based Activity Recognition and Interaction - WOAR '15*, ACM Press (2015), 1–11.
- [17] Møller, H., Sørensen, M.F., Jensen, C.B., and Hammershøi, D. Binaural Technique: Do We Need Individual Recordings? *Journal of the Audio Engineering Society* 44, 6 (1996), 451–469.
- [18] Myles, K. and Kalb, J.T. Vibrotactile Sensitivity of the Head. 2009.
- [19] Myles, K. and Kalb, J.T. Guidelines for Head Tactile Communication. 2010.
- [20] Myles, K., Kalb, J.T., Lowery, J., and Kattel, B.P. The effect of hair density on the coupling between the tactor and the skin of the human head. *Applied Ergonomics* 48, (2015), 177–185.
- [21] Paneels, S., Anastassova, M., Strachan, S., Van, S.P., Sivacoumarane, S., and Bolzmacher, C. What's around me Multi-actuator haptic feedback on the wrist. *2013 World Haptics Conference, WHC 2013*, (2013), 407–412.
- [22] Pielot, M., Krull, O., and Boll, S. Where is my team? Supporting Situation Awareness with Tactile Displays. *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10*, ACM Press (2010), 1705.
- [23] Precision Microdrives. Precision Microdrives 310-117 - Pico Vibe. <https://www.precisionmicrodrives.com/product/310-117-10mm-vibration-motor-3mm-type>.
- [24] Schneider, O.S., Israr, A., and MacLean, K.E. Tactile Animation by Direct Manipulation of Grid Displays. *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology - UIST '15*, ACM Press (2015), 21–30.
- [25] Scott, J.J. and Gray, R. A comparison of tactile, visual, and auditory warnings for rear-end collision prevention in simulated driving. *Human factors* 50,

- 2 (2008), 264–275.
- [26] Seo, J. and Choi, S. Initial study for creating linearly moving vibrotactile sensation on mobile device. *2010 IEEE Haptics Symposium*, IEEE (2010), 67–70.
- [27] Sloane, N.J.A., Hardin, R.H., Smith, W.D., and others. Spherical codes. URL <http://neilsloane.com/packings/>, (2000).
- [28] Tsukada, K. and Yasumura, M. *ActiveBelt: Belt-Type Wearable Tactile Display for Directional Navigation*. 2004.
- [29] Way, T.P. and Barner, K.E. Automatic visual to tactile translation. I. Human factors, access methods and image manipulation. *IEEE Transactions on Rehabilitation Engineering* 5, 1 (1997), 81–94.
- [30] Weber, B., Schätzle, S., Hulin, T., Preusche, C., and Deml, B. Evaluation of a vibrotactile feedback device for spatial guidance. *2011 IEEE World Haptics Conference, WHC 2011*, (2011), 349–354.
- [31] Zotkin, D.N., Duraiswami, R., and Davis, L.S. Rendering Localized Spatial Audio in a Virtual Auditory Space. *IEEE Transactions on Multimedia* 6, 4 (2004), 553–564.
- [32] Unity - Game Engine. <https://unity3d.com/>.