

Magnetips: Combining Fingertip Tracking and Haptic Feedback for Around-Device Interaction

Jess McIntosh

University of Copenhagen
Copenhagen, Denmark
jemc@di.ku.dk

Paul Strohmeier

University of Copenhagen
Copenhagen, Denmark
p.strohmeier@di.ku.dk

Jarrold Knibbe

Monash University
Melbourne, Australia
jarrod.knibbe@monash.edu

Sebastian Boring

University of Copenhagen
Copenhagen, Denmark
sebastian.boring@di.ku.dk

Kasper Hornbæk

University of Copenhagen
Copenhagen, Denmark
kash@di.ku.dk



Figure 1: Magnetips consists of (a) a magnetometer array to track the magnet on the fingernail, and (b) a coil to provide haptic feedback to it. These can be used simultaneously and integrated with mobile devices (c,d) to enable interactions in the space around the watch (e).

ABSTRACT

Around-device interaction methods expand the available interaction space for mobile devices; however, there is currently no way to simultaneously track a user's input and provide haptic feedback at the tracked point away from the device. We present Magnetips, a simple, mobile solution for around-device tracking and mid-air haptic feedback. Magnetips combines magnetic tracking and electromagnetic feedback that works regardless of visual occlusion, through most common materials, and at a size that allows for integration with mobile devices. We demonstrate: (1) high-frequency

around-device tracking and haptic feedback; (2) the accuracy and range of our tracking solution which corrects for the effects of geomagnetism, necessary for enabling mobile use; and (3) guidelines for maximising strength of haptic feedback, given a desired tracking frequency. We present technical and usability evaluations of our prototype, and demonstrate four example applications of its use.

CCS CONCEPTS

• **Human-centered computing** → **Interaction devices.**

KEYWORDS

Haptics, Mobile, Around-Device, Magnetic

ACM Reference Format:

Jess McIntosh, Paul Strohmeier, Jarrod Knibbe, Sebastian Boring, and Kasper Hornbæk. 2019. Magnetips: Combining Fingertip Tracking and Haptic Feedback for Around-Device Interaction. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019)*, May 4–9, 2019, Glasgow, Scotland Uk. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3290605.3300638>

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 2019, May 4–9, 2019, Glasgow, Scotland Uk

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300638>

1 INTRODUCTION

Research is exploring ways to extend interaction beyond the physical boundaries of our devices, through *Around-Device Interaction* (e.g., [5, 11, 19, 22, 29]). This interaction style is especially promising for mobile and wearable devices, as their small displays can limit the available space for interaction and suffer from fat-thumb occlusion [6].

Integrating these techniques with mobile devices presents a range of challenges. The tracking techniques, for example, variously suffer from occlusion (e.g., [29]), provide only coarse or two-dimensional positions [14], or require specific lighting requirements [19]. The feedback techniques either require active instrumentation of the user (e.g., [28, 34]), only work in limited directions (e.g., [36]), or have yet to be demonstrated in mobile form factors (e.g., [28, 36]).

We present Magnetips, a device that uses magnetism to enable tracking and haptic feedback for around-device interaction on mobile devices. This combination enables full 3D tracking and feedback above, around, and below the device, all within a small mobile form-factor. Magnetips is the first realisation of a combined technique for around mobile device interaction that works in any spatial direction and through visual occlusions. Magnetips presents the first example of passive magnetic tracking in 3D that accounts for the effects of the earth's magnetic field; affording a truly mobile setup with greater accuracy than previously demonstrated.

We describe (1) the implementation of Magnetips, (2) an evaluation of the tracking accuracy, both whilst the device is stationary and in motion, and (3) a psychophysics study of the user perception of the feedback, describing the signal parameters that generate the strongest haptic sensation, and the relationship between feedback strength and position.

These evaluations show that the maximum perceivable range of feedback is 56.6mm (at which distance tracking error is 6.38mm when under motion), and the ideal parameters (when tracking below 83Hz) are 12ms signals consisting of 4ms pulses. This allows a tracking frequency of up to 83Hz, beyond the frequency of most displays. We show that with geomagnetism cancellation algorithm, we improve the tracking accuracy by 17.4% while the device is under motion.

2 MAGNETIPS

Magnetips is a device that combines tracking and haptic feedback for around-device interaction. Figure 1 shows our device, consisting of: a copper coil, four small magnetometers, an IMU, a motor driver, a power source, and a magnet on the fingernail. The magnet (10x10x1mm disc magnet) is adhered to the users' fingernail. Magnetips can be integrated in a smart-phone case or snapped on to a smart-watch, to allow for retrofitting onto existing systems. We envision that the system may also be integrated within the device itself.

As the user moves their finger around the device, the four magnetometers collect field strength readings and estimate the position of the magnet in 3D. To create haptic feedback, we generate a magnetic field through the coil and alternate the polarity in order to induce vibration of the magnet. To reduce the size and weight requirements of the device, we use a coil without a magnetic core.

The high frequency of Magnetips allows feedback and tracking to be interleaved, whilst maintaining a high tracking rate (>60Hz, the tracking rate of most devices' screens). Another benefit of a magnet-based approach is that the tracking and feedback is not challenged by occlusion. Magnetips works through materials (such as sleeves, pockets, and bags), electronics (through displays and internal device circuitry), and through the body. Magnetips also works in three dimensions around the device, supporting above, beside, and behind device interaction.

We include an IMU to track the orientation of the device. This enables Magnetips to reduce the effect of geomagnetism and maintain accurate tracking in mobile scenarios. To our knowledge, this is the first work that uses the orientation of the device to compensate for errors when tracking a passive magnet in 3D.

As a result of the benefits of magnetic tracking and feedback, Magnetips affords a range of novel application scenarios. It allows tracking in all directions around the device. This enables behind device control and feedback when interacting with a smartphone, for example. This can also enable on-skin interaction and feedback beside the device. As magnetic input and output techniques are not effected by occlusion from most common materials, Magnetips affords eyes-free interaction through clothes. We discuss the technical implementation and example applications in detail later.

3 RELATED WORK

There have been a variety of approaches to using the space around mobile devices for interaction. Capacitive sensing has been used to extend the input abilities of smartwatches [22] and extend the sensing area of smartphones to include the area above a phone [15]. Infrared (IR) depth-sensors have been used to track regions off the device, within a line of sight of the sensors [7, 19]. Similarly, computer vision based systems have been deployed in wearable technology [29]. Alternative approaches include using the skin as an interactive area [23] or physically extending the size of mobile devices, by, for example, extending the interaction space to a watch's wristband [24, 31] or to integrate with clothing [26]. Recently, electric field sensing has enabled 3D interaction around mobile devices [20, 37]. Although there is much work on increasing the input space, there is little complementary work on providing haptic feedback at the tracked point.

Mid-air Haptics

The mobility and size constraints of wearables and smartphones present challenges when integrating a mid-air haptic feedback system. Vibration, the most common feedback modality in mobile devices, for example, requires surface contact. To be used for off-device feedback, a vibration motor needs to be worn by the user (e.g., [11]), which may impede natural hand interactions.

Several mid-air haptic feedback approaches do not require active components to be worn on the hands. Ultrahaptics uses focused ultrasound to create haptic feedback in mid-air [8]. Jets of air [33] and air vortices [27] have also been explored. While these approaches work well for larger, stationary applications, they currently lack portability, which renders them unsuitable for integration with mobile devices. In addition, they require an un-occluded position on the surface of the device to function. In contrast, magnetism can work through materials such as glass, ceramics, plastics and non-ferromagnetic metals, as commonly used in devices.

Spelmezan et al. presented a method that provides haptic feedback using electrical arcs [28]. However, this technology has only been demonstrated to work close to the surface of the device, which enables hover interaction up to 4mm.

Finally, in FingerFlux, Weiss et al. showed that by attaching a permanent magnet to the fingertip, attracting and repelling forces can be felt via electromagnets [35]. By using an array of electromagnets, positioned beneath a surface, Weiss et al. were able to guide a users' finger during screen-based interaction. We build on this work and tailor it to a mobile solution. We replace the array of small electromagnets with a single large coil. This design allows us to reduce the size, weight and power requirements of the FingerFlux approach.

A benefit of using a passive magnet haptic feedback is that we can simultaneously use magnetic approaches to track the users' finger around the device. This tracking benefits from the same features as the haptic system: tracking in three dimensions and through materials. In the next section, we present the existing literature on magnetic tracking.

Magnetic Tracking

There is much work in the area of magnetic sensing for magnet position estimation. In Finexus [10], users wore electromagnets on their fingertips to enable accurate tracking. We instead envision a system that relies on only passive instrumentation of the user, affording a less obtrusive and untethered setup. Therefore, we will focus on passive magnet tracking.

Research has explored the use of both single magnetometers, for lightweight yet spatially-restricted tracking, and multiple magnetometers, for full 3D tracking.

As an example of single magnetometer use, Ketabdar et al. explored gesture tracking around the device while the user wears/holds a permanent magnet [18]. They used the euclidean norm of the magnetic field strength in order to determine the magnets' distance from the device. While this did not enable precise location tracking, it did support gesture recognition proximal to the device. In similar work, Ashbrook et al. [1] devised a ring with an embedded magnet, where rotations and translations of the ring could be used as an additional input modality. Radial movements around a device can also be sensed by a single magnetometer as Harrison and Hudson demonstrated in *Abacadabra* [14].

In a variation on the device-mounted magnetometer approach, Han et al. [13] instrumented the users' wrist with two magnetometers, in order to support 2D mid-air handwriting. In *uTrack*, Chen et al. [9] demonstrated 3D tracking of magnets with as few as two magnetometers.

One of the challenges with using magnetometers to track magnets is the effect of the Earth's magnetic field, which needs to be taken into account during calibration, yet changes with any change in sensor orientation [10]. This can result in significant declines in tracking accuracy as a magnet moves just a few centimeters away from the sensors [10]. There have been a number of approaches to dealing with this challenge, from restricting input to limited dimensions [13], using large magnets in close proximity to the sensors [9], to using advanced signal filtering with electromagnets [10].

We present another novel solution to this approach, by using an IMU (inertial measurement unit) in order to track changes in orientation, which can then be used to cancel the effects of geomagnetism. This allows us to maintain accurate tracking without any limit to interaction, without a strong magnet, and without the use of electromagnets.

4 IMPLEMENTATION

Magnetips consists of two elements, a tracking system and a haptics system. We will describe this in the following.

Hardware

Tracking. We use four three-axis magnetometers (LIS3DML), placed in a rectangular arrangement. Figure 2 shows the arrangement of the sensors, and the exact placements of the sensor ICs within the board (32.5 x 26.7 mm rectangle, small enough to fit within the dimensions of an Apple watch series 3, see Figure 1). Additionally, we use an IMU (LSM6DS33) to track the orientation of the sensor board / device. The sensors are able to sample at a frequency of 1kHz. This high frequency allows us to multiplex the haptics with the tracking, and is one of the main reasons for choosing this sensor.

Haptics. We provide haptic feedback by creating electromagnetic pulse-bursts with a hand-wound coil. As we intended




Figure 2: The array of sensors consisting of 4 magnetometers (red) and an accelerometer and gyroscope (blue).

to use this technology with mobile devices, using the coil without a magnetic core is important to reduce the weight and size of the device. The coil is 45 mm by 55 mm and has 38 windings of 0.5 mm enamelled copper wire. The coil is driven with a Pololu VNH5019 Motor Driver Carrier. Each burst consists of a single pulse or a series of pulses with varying frequencies.

Controller. Haptic feedback and sensors measurements are controlled by a Teensy 3.6. The measurements are sent to a desktop PC via serial communication, ready to be processed by a host application.

Software

Tracking Algorithm. We use the same algorithm as Chen et al. in Finexus [10] to estimate the position of the magnet in 3D. This algorithm uses the magnetic field strength that each magnetometer is subject to, in order to estimate the distance between the magnet and each magnetometer by using the inverse cubic relationship of field strength to distance.

Geomagnetism cancellation. The earth’s magnetic field produces a bias in the magnetometers readings. The field consists of an inclination (the deviation between true north and magnetic north), a declination (the angle between the magnetic field lines and the earth’s surface) and an intensity. These three components can be represented as transformation matrix M , which describes the rotation between a true north coordinate system and the earth’s magnetic coordinate system. The field is relatively constant throughout a small geographical region, say a city, and can be calculated based on latitude and longitude ¹.

Geomagnetism, however, is a hindrance to the tracking of magnetic objects as it skews the readings of magnetometers. During interaction, the wrist of the user will move (even if slightly), thus altering the readings of the magnetometers. Simultaneously, the magnet’s effect on the sensors will also

be present. It then becomes challenging to know how much signal change can be attributed to movement of the magnet or movement of the device.

Our algorithm is one of the ways to overcome this challenge. To begin with, the sensor’s orientation (in world coordinates) has to be known. Yet, a magnetometer by itself cannot estimate the orientation of a device with 3 DoF (degrees of freedom) accurately. In the same way, an accelerometer cannot track the orientation with 3 DoF without drifting around the axis of gravity. A gyroscope can also be fused with the existing data, in order to give a more responsive estimation. Each of these sensors has 3 axes, for a total of 9 axes. A system that performs such an orientation estimation is called an attitude and heading reference system (AHRS). We implement Madgwick’s AHRS to get accurate estimations of the orientation in 3 DoF.

However, in the presence of a magnet, 9-axis AHRS cannot be used due to the magnet’s influence on the magnetometer. Fortunately, 6-axis AHRS can be used which uses only the accelerometer and gyro [21], which is sufficient for short interaction times. Over time, however, there will be accumulated drift due to the absence of magnetometer data. This minor accumulated drift can be corrected for after the magnet is taken away when the interaction is complete. The system also needs to know when to switch from 9-axis to 6-axis AHRS. We implement this detection with a simple threshold on the difference between the magnitude of the earth’s magnetic field vector and the current measurements. The AHRS is likely to have been given some biased magnetometer data shortly before reaching this threshold. To avoid this issue, we store the IMU data into a buffer, then recalculate the 6-axis AHRS from the past second of data.

The key to our geomagnetism cancellation algorithm is to reverse the AHRS process in order to estimate magnetometer data from the orientation data. That is, given an absolute orientation of the device, we can calculate the expected magnetometer readings (using the local earth’s magnetic field components, represented as transformation M and intensity I_M) for each axis of the magnetometers. We do this as follows for each magnetometer: first, we calculate the representation of a magnetometer’s axes (as vectors) in the earth’s magnetic coordinate system using the earth’s magnetic field transform M , the intensity I_M and the sensor orientation S_R . Second, we project these transformed vectors onto the earth’s magnetic field vector (the x -axis of M). The calculations are:

$$\begin{aligned} I_{S_x} &= M^{-1} \cdot S_R \cdot (1, 0, 0)^T \cdot (I_M, 0, 0)^T \\ I_{S_y} &= M^{-1} \cdot S_R \cdot (0, 1, 0)^T \cdot (I_M, 0, 0)^T \\ I_{S_z} &= M^{-1} \cdot S_R \cdot (0, 0, 1)^T \cdot (I_M, 0, 0)^T \end{aligned}$$

Resulting values fall into $[-I_M, +I_M]$. For example, if the x -axis of a magnetometer is properly aligned with the earth’s

¹Magnetic Field Calculator: <https://www.ngdc.noaa.gov/geomag-web/>



Figure 3: A setup of how the magnet and sensor board was tracked using OptiTrack.

magnetic field, it will have the full intensity, while both the y -axis and z -axis will read 0. The actual magnetometer readings are then corrected using the expected readings $I_{S_x, y, z}$. In an ideal case (i.e., no magnetic object being present near the sensor), the corrected readings would be 0 on all axes.

5 TRACKING EVALUATION

Next, we evaluate the accuracy of our tracking solution. Importantly, we demonstrate the effectiveness of our novel geomagnetism cancellation, showing an improvement of 17.4% tracking accuracy when the device is under motion. We also show how the magnetic field produced by our haptics system might interfere, and show that the minimum delay needed before and after producing a haptic signal is 2ms.

Tracking Accuracy

Static sensors. We tested the tracking accuracy by first keeping the sensor board static and moving the magnet in a 3D volume, centred around the sensors. For this condition, we calibrated the geomagnetism before measuring the data. As the sensor base is static, our geomagnetism cancellation algorithm does not provide any further benefits. In each of these tests, Magnetips was placed on a table. The magnet was attached to retro-reflective markers for optical tracking. We moved the magnet-marker in an $80 \times 80 \times 80$ mm volume and tracked its position with both Magnetips and OptiTrack, where OptiTrack provided the ground truth for this data (Figure 3). Our results are shown in table 1. We present these results with varying volume sizes. The error is calculated by the euclidean distance between the OptiTrack position and the Magnetips position (mm) for each frame of the data. The error for each frame is then averaged. The magnetometer data sampling rate was set to 1KHz, with ± 8 gauss sensitivity. The results show that there is an approximately linear increase in error from the 40 mm to the 80 mm cube lengths.

Moving sensors. The previous results, however, show the upper bound of tracking accuracy where the sensing base




Figure 4: A plot of the average magnetometer errors vs the time since a haptic pulse had ended.

is not moving and therefore the geomagnetism is constant. Therefore, we conducted two further tests. Firstly, we fixed the magnet at a known position away from the sensor board, using the green cage shown in figure 3 (in the middle indentation). We then recorded the data whilst moving and tilting the sensor board and magnet. In this test, the previously used algorithm performs poorly and results in an average error of 7.790 mm. After using our geomagnetism cancellation technique based on the AHRS, our findings show an average error of 5.514 mm (see table 1 for comparisons between geomagnetism cancellation on and off). This is an accuracy improvement of 2.276 mm, for a magnet position fixed at 5cm above the sensor base.

Then finally, to simulate a more realistic scenario, we moved the base and the magnet simultaneously. Again, with our geomagnetism cancellation we observed reductions in the tracking error in each of the conditions. The mitigation technique seems to have a greater effect with a larger volume, probably due to the fact that the earths magnetic field has a stronger effect relative to the effect of the magnet when further away. In the largest volume, we see a decrease in the error of the tracking by 1.585 mm.

GC	Sensor base	Magnet	Length of cube (mm)		
			40	60	80
N/A	Static	Moving	3.001	5.055	6.686
Off	Moving	Static*	7.790		
On	Moving	Static*	5.514		
Off	Moving	Moving	4.408	7.502	9.095
On	Moving	Moving	4.354	6.379	7.510

Table 1: The average errors accumulated over the duration of data sampling in mm. Geomagnetism cancellation abbreviated to GC. *Relative to the sensor base.

Multiplexing Tracking With Haptics

The haptics system we use inevitably interferes with the magnetometer readings without careful timing. We investigated whether there is a lasting effect from the coil after the controller turns off power to it, or whether the latency of the reading is an issue. Indeed, we found that if we begin to sample data just afterwards, there is some error in the readings as shown in Figure 4. The error in this graph represents the average error from the earth's magnetic field on a logarithmic scale. The interference stops after a 2ms delay, but it is very large and unpredictable before this. Due to a lack of control over timing with the hardware (continuous 1KHz sampling), we also added a 2ms delay after reading and before sending a haptic signal. With bespoke hardware to improve the timing of the sampling, these delays can be reduced significantly. In total, there must be a 4ms delay added onto the sampling period. This limits the sampling frequency of the tracking if haptic pulses are to be used often during the interactions. We discuss the implications of this in the next section.

6 HAPTICS EVALUATION

We conducted a series of evaluations to demonstrate haptic feedback properties. Tracking accuracy can simply be measured, but the experience of haptic feedback is subjective and not as easy to quantify. Therefore, we spend more time on the feedback evaluation than on the tracking, though both parts are equally important for Magnetips.

Although there is past work that explores vibrational feedback with similar parameters using conventional vibrotactile stimulation, there are a number of factors which could influence the transfer from signal frequency to experienced vibration. The main two factors include: the location of the




Figure 5: In both experiments, the participants finger was fixed in a single position. Top left shows the setup for the first experiment. Towards the bottom, the apparatus can be seen for moving the coil for the second experiment.

magnet on the fingernail; the transfer of electrical to electromagnetic to mechanical energy.

The haptics evaluation consists of three parts. First, to better understand how to design haptic signals, we a) present a user study of possible parameters, given constraints in tracking frequency. Using the parameter combination that lead to the strongest feedback in study 'a', we then b) present a user study that investigates how the finger's position relative to the coil affects the strength of the feedback. Finally, to better understand variations in strength discovered in study 'b', we c) present measures of the electromagnetic field produced when providing feedback.

a) Feedback Parameters

We conducted an experiment to better understand what parameter combination provides the strongest haptic sensation. We recruited 10 participants. We affixed a magnet to the fingernail of the index finger of the dominant hand of each participant with double sided adhesive pads. For this experiment, we fixed the position of the finger (and thus magnet) with respect to the coil as illustrated in Figure 5, inlay. Participants were asked to rate how strongly they perceived haptic signals which varied in duration and frequency.

Independent Variables

Signal Duration. The maximum duration of each signal is limited by the desired tracking frequency. As the electromagnetic field required for providing haptic feedback interferes with the magnetic tracking, tracking and feedback must be alternated. We found that we need to delay readings by 3ms after generating the electromagnetic field to ensure error free sampling of the magnet position. Measuring the magnets position takes 1 ms. Tracking at frequencies common with current touch solutions (60Hz or above) provides us with a total time window of at most 16.7 ms, which supports a maximum feedback duration of 12.7 ms (per the results of the section 5).

Pulse Length. The perception of a haptic signal is strongly linked to its frequency content [4]. To explore how the frequency effects perceived strength, we vary the length of electromagnetic pulses that each signal consists of. Signal Duration and Pulse Length interact: the Pulse Length cannot be longer than the Signal Duration, short Pulse Length enable signals consisting of multiple pulses, which might create stronger haptic signals. The Pulse Length is the reciprocal of its frequency.

Combinations. To better understand the parameter space, we tested 11 combinations of Signal Duration and Pulse Length (12 ms duration: 2, 3, 4, 6 & 12 ms pulse length, 6ms duration: 2, 3 & 6 ms pulse length, 4 ms duration: 2 & 4ms pulse length, 2 ms duration: 2 ms pulse length)




Figure 6: Perceived strength of feedback for: a) All combinations of Duration and Pulse Length, b) Duration, c) Pulse Length.

Dependant Variable & Procedure

For each combination, participants rated how strongly they perceived the haptic feedback. Following methodology suggested by Gescheider [12], we asked participants to freely assign values to the feedback strength. These ratings not only provide us with a guideline on choosing a Signal Duration / Pulse Length combination, but also allow us to create visual response scales for each independent variable, as done by Strohmeier and Hornbaek [32]. Participants rated each combination three times. Combination order was randomised.

Results

User ratings were normalized so that the maximum value per user was 1 and ratings were then averaged per user and condition using the geometric mean. We subtracted each users average rating from all values, so as to move all users data in the same frame of reference, where zero represents the average rating, positive values are above average and negative values below average, adapted from Strohmeier and Hornbaek [32]. The resulting scales are unit free, but the confidence intervals shown in Figure 6 a,b&c provide guidance concerning size of effects.

Overall the combination of 12ms signal duration with 4ms pulses yielded the highest perceived strength. This result agrees with the literature, 4ms pulses produce a 250Hz signal, which is in the centre of the frequency range Pacinian cells respond to [16], stimulating the same cells responsible for texture perception [4]. The combination of 12ms signal duration and 4ms pulse duration supports tracking speeds up to 83Hz. If the tracking speed is increased to 166Hz or above (6ms duration or above) the strength of feedback that can be provided with a discrete signal becomes significantly weaker. Mean results and corresponding 95% confidence intervals for all combinations are shown in Figure 6a.

To better understand the effects of Duration and Pulse Length on their own, we also plot them individually. Figure 6 b shows what appears to be a relatively linear relation between the signal's length and it's perceived strength for signal durations <12ms. Figure 6 c shows that the intensity

levels we found match those recorded using direct nerve readings of pacinian cells [25]. It should be noted that not all pulse lengths are combined with all signal durations, the graphs only show the combinations described previously. This is reflected in the wider confidence intervals for short durations and for the 4ms and 12ms pulses.

b) Actuation Volume

The second experiment was designed to explore the uniformity of the haptic feedback based on where the tracked point is positioned relative to the coil, as well as establish an ideal range within which to provide haptic feedback. Based on the result of the previous study, we conducted our exploration of the actuation volume using a feedback duration of 12ms with a pulse duration of 4ms.

Apparatus. The apparatus for this experiment can be seen in figure 5, which uses a modified 3D printer, that moves the coil in 3 dimensions. The coil is held above the printer by acrylic to avoid magnetic field interference that may be present nearer to the bed of the printer. As with the previous experiment, the participant is asked to keep their finger in a fixed location throughout. The experiment moved the coil into 60 different positions around the device, in a 60x60x60 mm cube. The area inside the coil was not sampled.

Figure 7 illustrates the volume that was sampled, from a Z height that starts with the coil level with the fingernail. At each location the participant was asked to rate how strongly they perceived the strength of the feedback. During the experiment, the participants were blind-folded and wore headphones to avoid any visual or audial biases. Each of the 60 locations were repeated 3 times.

Since the field strength is symmetric about 3 axes, we decided to only sample an eighth of the field. The data is extrapolated in Figure 7 to mirror the x and y axes.

Variables

The independent variable of this experiment is the position of the finger in 3D space. The origin of this position is in the

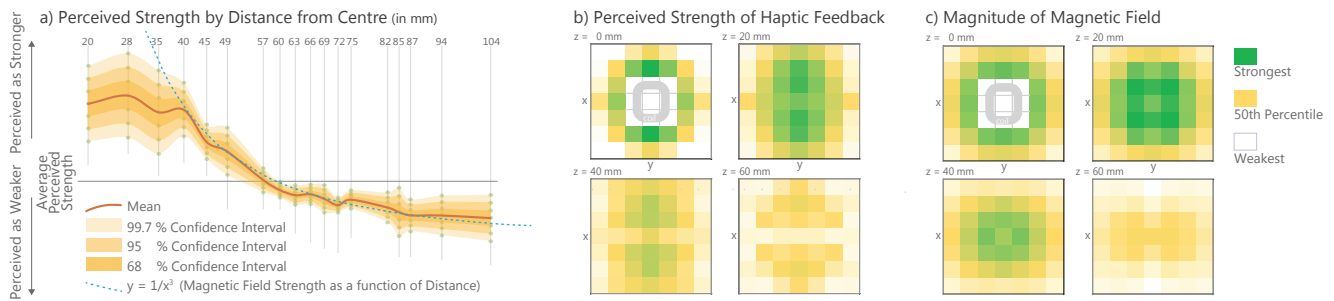


Figure 7: a) Perceived Strength of Feedback sorted by distance to centre. b) Strength of Haptic Feedback as rated by participants. c) The magnitude of the magnetic field as measured with a magnetometer

centre of the coil. We use the euclidean distance from the fingertip to the centre of the coil as a measurement to compare the strength against. Our dependant variable is the perceived strength of the feedback, as in previous experiment.

Results

We processed the strengths ratings the same way as we did for the previous experiment. We found that, up to 4cm, the ratings were relatively similar. Beyond that, the mean results aligned with $strength = 1/distance^3$ ($R = .993$), which is the behavior one might expect if the feedback strength were directly proportional to magnetic field strength. Figure 7, left, shows the perceived strength plotted against the euclidean distance of the location to the centre of the coil.

To understand the range that can reasonably be used for haptic feedback, we analyzed the differences in strength ratings along the distance axis. We assume that, moving from centre of the coil outward, users can clearly feel the haptic signal, as long as there is a significant difference to a weaker signal further out. We found that 56.6 mm was the last value to be significantly different from weaker signals further out (Bonferroni corrected, $p < 0.01$ for 84.9, 87.2, 93.8 & 103.9). While individual users might experience the signal beyond that, 56.6mm is the furthest out that users could clearly distinguish the feedback from even weaker signals. The full confusion matrix can be found in the supplementary material.

To better understand how the positioning of the finger relative to the magnetic field effected the ratings, we created a heat-map figure 7, centre. The heat map demonstrates that, while the strength generally declines with distance, the pattern is more complex. For example, at $z = 60$ mm the area directly on top of the centre of the coil provides relatively weak feedback, compared above or below the centre on the x-axis. With knowledge of such patterns, haptic widgets can be created that are even further away from the device than 56.6mm, if one places them in areas where the feedback can

be clearly perceived. Alternatively, these variances in perceived strength could be corrected for, by taking the position of the finger into account when providing feedback.

c) Measurements of Magnetic Field

If the perceived strength correlated directly with the magnetic field, we could leverage the existing know-how of magnetic fields for strength calibration or widget-placement strategies, as discussed above. To better understand the patterns found in the previous user study, we therefore repeated the procedure. This time, instead of asking participants to rate the strength of the vibration, we placed a magnetometer where the participants finger would be and measured the strength of the magnetic field produced by the coil.

Results

The magnitude of the field strength correlates well with the haptic feedback ($R = .805$, Figure 7, right) but does not completely explain all variations. At $z = 60$ mm, the magnitude hardly correlated with the perceived strength at all ($R = -0.01$). Instead the perceived strength correlated well with the strength measured with the y-axis of the magnetometer ($R = .84$). Looking at the measured dimensions of the magnetic field individually demonstrates that they contribute to the haptic experience differently: the y-dimension of the magnetic field measures appeared most strongly to influence perception ($R = .721$), followed by z ($R = .62$). The x dimension did not correlate well with perceived strength of the feedback ($R = .08$). Plots for the individual dimensions can be found in the supplementary material.

Summary of Findings

Users could feel signals best at a frequency between 250 to 333Hz. Within the short durations tested, there was a linear relationship between perceived strength and signal duration. The longest possible signal duration (12ms) pairs nicely with 4ms pulses (250Hz), which was the combination users rated as strongest. Using the 12ms, 250Hz combination,



Figure 8: Magnetips enables these interactions: a) 3D mid-air, b) Behind arm, c) Back of device, d) Through material.

we explored the volume around the coil and found that users could reliably detect the feedback to at least 56.6mm. We also found that there were non-linear effects, based on the relative position of the finger to the centre of the coil. This pattern correlated with the magnetic field strength to some extent ($R = .84$), but the exact relation requires further exploration.

7 APPLICATIONS

Magnetips can track and provide feedback in three dimensions around the device, and works regardless of visual occlusions. This enables a range of novel interaction scenarios. To demonstrate some of these opportunities, we built example applications that uses Magnetips capabilities, as shown in figure 8 and in the video figure. All the applications that we have built use haptic feedback in mid-air, which is the main benefit that Magnetips brings to around-device interaction.

3D Mid-air Interaction

In past literature, there has been work that demonstrates tracking of permanent magnets even in 3 dimensions. However, by adding haptics, this enables virtual elements in mid air to be feelable.

To demonstrate 3-dimensional interaction, we created a clock application (figure 8 a). This application lets the user choose the time by rotating the finger in a radial movement around the device, similar to Abracadabra [14]. In this version, we let the user choose between minutes and hours by varying the depth of the radial movement. In addition, each step through a unit of time (second or minute) triggers a haptic signal. In this example, we demonstrate above device tracking, side of device tracking, and mid-air feedback.

Behind Arm

Magnetips enables interactions through occlusions made by the body itself. We prototyped an application for panning a map in 2D, by dragging the map underneath the arm (figure 8 b). Such an interaction modality may be useful when visual attention is still required for the task, as this does not occlude the display of the wearable device [3]. This interaction is also

an example of skin input, which Magnetips can complement with localised haptic feedback.

Back of Device

Back of device interaction is a common research topic in mobile device interaction [2], for smartphones in particular. We built a photo editing application to demonstrate that Magnetips also enables this interaction paradigm. In the application, while the user is preparing a photo for sharing on their smartphone, the user could press and hold the 'filters' button and then use their index finger behind-the-device to scroll through the available list, receiving haptic feedback at every item boundary. This allows the user to also get visual feedback of the effect of the filter, without losing any on-screen space to displaying the list. Similarly, the technique could be used to create mid-air triggers with feedback behind the device for gaming.

Eyes-free Interaction & Interacting Through Materials

Magnetips can support eyes-free interaction in different ways. Current devices typically use vibration as eyes-free feedback channel. This causes the whole device to vibrate and can cause a buzzing sound when in contact with surfaces. Using Magnetips, the user can receive feedback off-device, directly onto their finger - supporting subtle interaction. A similar interaction can also work through clothes, for example, through a sleeve for a smartwatch or through trouser pockets, bags and purses for a smartphone.

8 LIMITATIONS & FUTURE WORK

Magnetips enables reliable haptic sensations up to 56mm from the center of the coil. At that range, the tracking error is 6.38mm. This enables Magnetips to create a larger interaction space around mobile devices. Magnetips presents a range of new opportunities, which we explored above with our applications. There are, however, also limitations to Magnetips, which remain future work.

Coil Design and Power Usage

We designed Magnetips to use a large coil that can easily in-case existing devices. In part, this was inspired by the Qi charging coils already found in many devices.

We used a single coil for providing feedback around devices. In future work, different sizes of coils, tessellations of coils, and their impact on the haptic experience could be explored. For example, a phone-sized Magnetips device could consist of one large coil, spanning the majority of its size, or multiple smaller coils. By creating different designs of coil arrays, designers could experiment with different granularities of feedback in certain locations around the device. For example, an additional small coil at the bottom of a device could be used to drive a stronger sensation above the location of the traditional 'home' button.

While Magnetips has been designed to fit within mobile devices, we have yet to optimise its power requirements. The power supply in the studies was set to use a maximum of $20V \times 4A = 80W$, for 12 ms per tick of feedback. This equates to less than 1 Joule of energy (0.9576 Joules). This is much more than current vibration motors, for example, which would require ~ 0.02 Joules for the equivalent feedback. (For reference, the iPhone 7 can store 39,600 joules). The power consumption may be reduced by using an array of smaller coils. As the position of the magnet is known, one can then generate a directed magnetic field towards the position of the magnet, similar to the electromagnet array in FingerFlux [35]. This would allow the same strength of haptic feedback at a fraction of the power.

In the studies we used a power supply rather than battery to remove any variation in energy use due to the battery discharging over the course of the study. However, a mobile version of the haptics system is feasible, with a 7.4V 600mAh 25C battery, as shown in the video figure and in Figure 9.

Magnetips is limited to the use of a single magnet. Passive magnets cannot be individually enabled and disabled, therefore it remains a challenge to address individual magnets for




Figure 9: The haptics system can be operated with very few components, requiring only a coil, motor driver, battery and a microcontroller.

either tracking or actuation. The tracking of multiple magnets may be solvable with a more complex optimiser and with more magnetometers, but we are currently unaware of an existing solution. However, for haptic feedback, a coil array may provide coarse addressability at the very least.

Magnet Shape

Users are required to wear a magnet. While this can be considered a drawback, MEMS magnetometer sensors are becoming increasingly better in sensitivity. This means that with future improvements to technology, smaller magnets can be used to achieve the same level of tracking performance, making it easier and more viable to embed them into nail art [17], or even into the finger [30]. Aside from wearing the magnet, it is possible to embed magnets into interactive tools, such as a stylus, to track and provide feedback in mid-air.

9 CONCLUSION

Magnetips improves on previous work by not only preventing occlusion of small displays, but also allowing users to interact whilst glancing away from the device, as users can receive haptic guidance. Magnetips does not require line-of-sight between the device and the tracked point, allowing for interaction in locations previously not possible with haptic feedback. This enables, for example, on-skin input and output behind the arm.

Through a series of technological evaluations and user studies, we showed that the current implementation can track the users finger with an average error of $< 6.4mm$ during a mobile task within a volume of $60 \times 60 \times 60mm$. Our results show that users can reliably feel the haptic feedback up to at least 56mm. We show that we can track the users finger at 60Hz and provide concurrent haptic feedback that is clearly perceivable. We suggest that haptic signals are optimised for strength at 12ms duration and 250Hz. If higher tracking and feedback frequencies are required, we suggest using 6ms duration at 330Hz for up to 166Hz tracking.

We presented four basic usage scenarios - aimed to demonstrate how magnetips might benefit familiar interactions. We believe, however, that the opportunities offered by the unique combination of colocated tracking and haptic feedback, in combination with the ability to interact through occluding materials, might extend far beyond these scenarios, and look forward to seeing how magnetips might be adapted in other form factors and contexts.

10 ACKNOWLEDGEMENTS

This work was supported by the European Research Council under the European Union's Horizon 2020 research and innovation program (grant agreement 648785).

REFERENCES

- [1] Daniel Ashbrook, Patrick Baudisch, and Sean White. 2011. NENYA: subtle and eyes-free mobile input with a magnetically-tracked finger ring. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2043–2046.
- [2] Matthias Baldauf, Sebastian Zambanini, Peter Fröhlich, and Peter Reichl. 2011. Markerless Visual Fingertip Detection for Natural Mobile Device Interaction. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '11)*. ACM, New York, NY, USA, 539–544. <https://doi.org/10.1145/2037373.2037457>
- [3] Patrick Baudisch and Gerry Chu. 2009. Back-of-device Interaction Allows Creating Very Small Touch Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 1923–1932. <https://doi.org/10.1145/1518701.1518995>
- [4] Sliman J Bensmaia and Mark Hollins. 2003. The vibrations of texture. *Somatosensory and Motor Research* 20, 1 (2003), 33–43. <https://doi.org/10.1080/0899022031000083825>
- [5] Joanna Bergstrom-Lehtovirta, Sebastian Boring, and Kasper Hornbæk. 2017. Placing and Recalling Virtual Items on the Skin. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*. ACM Press, New York, New York, USA, 1497–1507. <https://doi.org/10.1145/3025453.3026030>
- [6] Sebastian Boring, David Ledo, Xiang 'Anthony' Chen, Nicolai Marquardt, Anthony Tang, and Saul Greenberg. 2012. The Fat Thumb: Using the Thumb's Contact Size for Single-handed Mobile Interaction. In *Proceedings of the 14th International Conference on Human-computer Interaction with Mobile Devices and Services (MobileHCI '12)*. ACM, New York, NY, USA, 39–48. <https://doi.org/10.1145/2371574.2371582>
- [7] Alex Butler, Shahram Izadi, and Steve Hodges. 2008. SideSight: Multi-ÅtouchÅ Interaction Around Small Devices. In *Proceedings of the 21st annual ACM symposium on User interface software and technology - UIST '08*. ACM Press, New York, New York, USA, 201. <https://doi.org/10.1145/1449715.1449746>
- [8] Tom Carter, Sue Ann Seah, Benjamin Long, Bruce Drinkwater, and Sriram Subramanian. 2013. UltraHaptics: multi-point mid-air haptic feedback for touch surfaces. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*. ACM, 505–514.
- [9] Ke-Yu Chen, Kent Lyons, Sean White, and Shwetak Patel. 2013. uTrack: 3D input using two magnetic sensors. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*. ACM, 237–244.
- [10] Ke-Yu Chen, Shwetak N Patel, and Sean Keller. 2016. Finexus: Tracking precise motions of multiple fingertips using magnetic sensing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 1504–1514.
- [11] Euan Freeman, Stephen Brewster, and Vuokko Lantz. 2014. Tactile Feedback for Above-Device Gesture Interfaces: Adding Touch to Touchless Interactions. In *Proceedings of the 16th International Conference on Multimodal Interaction (ICMI '14)*. ACM, New York, NY, USA, 419–426. <https://doi.org/10.1145/2663204.2663280>
- [12] George A Gescheider. 1988. Psychophysical Scaling. *Ann. Rev. Psychol* 39 (1988), 169–200.
- [13] Xinying Han, Hiroaki Seki, and M Hikizu. 2007. Wearable handwriting input device using magnetic field. In *SICE, 2007 Annual Conference*. IEEE, 365–368.
- [14] Chris Harrison and Scott E. Hudson. 2009. Abracadabra. *Proceedings of the 22nd annual ACM symposium on User interface software and technology - UIST '09* (2009), 121. <https://doi.org/10.1145/1622176.1622199>
- [15] Ken Hinckley, Seongkook Heo, Michel Pahud, Christian Holz, Hrvoje Benko, Abigail Sellen, Richard Banks, Kenton O Hara, Gavin Smyth, and Bill Buxton. 2016. Pre-Touch Sensing for Mobile Interaction. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016), 2869–2881. <https://doi.org/10.1145/2858036.2858095> arXiv:1605.09782
- [16] Roland S Johansson and J Randall Flanagan. 2009. Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature reviews. Neuroscience* 10, 5 (2009), 345–59. <https://doi.org/10.1038/nrn2621>
- [17] Azusa Kadamura and Itiro Siio. 2015. MagNail: user interaction with smart device through magnet attached to fingernail. In *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers*. ACM, 309–312.
- [18] Hamed Ketabdar, Mehran Roshandel, and Kamer Ali Yüksel. 2010. Towards using embedded magnetic field sensor for around mobile device 3D interaction. In *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services*. ACM, 153–156.
- [19] Jarrod Knibbe, Diego Martinez Plasencia, Christopher Bainbridge, Chee-Kin Chan, Jiawei Wu, Thomas Cable, Hassan Munir, and David Coyle. 2014. Extending interaction for smart watches. In *Proceedings of the extended abstracts of the 32nd annual ACM conference on Human factors in computing systems - CHI EA '14*. ACM Press, New York, New York, USA, 1891–1896. <https://doi.org/10.1145/2559206.2581315>
- [20] Mathieu Le Goc, Stuart Taylor, Shahram Izadi, and Cem Keskin. 2014. A low-cost transparent electric field sensor for 3d interaction on mobile devices. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 3167–3170.
- [21] Sebastian OH Madgwick, Andrew JL Harrison, and Ravi Vaidyanathan. 2011. Estimation of IMU and MARG orientation using a gradient descent algorithm. In *Rehabilitation Robotics (ICORR), 2011 IEEE International Conference on*. IEEE, 1–7.
- [22] Ian Oakley and Doyoung Lee. 2014. Interaction on the Edge: Offset Sensing for Small Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 169–178. <https://doi.org/10.1145/2556288.2557138>
- [23] Masa Ogata and Michita Imai. 2015. SkinWatch: Skin Gesture Interaction for Smart Watch. *Proceedings of the 6th Augmented Human International Conference on - AH '15* (2015), 21–24. <https://doi.org/10.1145/2735711.2735830>
- [24] Simon T Perrault, Eric Lecolinet, James Eagan, and Yves Guiard. 2013. WatchIt : Simple Gestures and Eyes-free Interaction for Wristwatches and Bracelets. In *Proc. CHI*. 1451–1460.
- [25] Robert F. Schmidt. 1986. *Fundamentals of Sensory Physiology*. Vol. 3. Berlin Heidelberg New York Tokyo. 25–33 pages. <https://doi.org/10.1017/CBO9781107415324.004> arXiv:arXiv:1011.1669v3
- [26] Stefan Schneegass and Alexandra Voit. 2016. GestureSleeve: using touch sensitive fabrics for gestural input on the forearm for controlling smartwatches. *Proceedings of the 2016 ACM International Symposium on Wearable Computers - ISWC '16* (2016), 108–115. <https://doi.org/10.1145/2971763.2971797>
- [27] Rajinder Sodhi, Ivan Poupyrev, Matthew Glisson, and Ali Israr. 2013. AIREAL: interactive tactile experiences in free air. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 134.
- [28] Daniel Spelmezan, Deepak Ranjan Sahoo, and Sriram Subramanian. 2017. Sparkle: hover feedback with touchable electric arcs. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 3705–3717.

- [29] Srinath Sridhar, Anders Markussen, Antti Oulasvirta, Christian Theobalt, and Sebastian Boring. 2017. WatchSense: On- and Above-Skin Input Sensing through a Wearable Depth Sensor Srinath. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*. ACM Press, New York, New York, USA, 3891–3902. <https://doi.org/10.1145/3025453.3026005>
- [30] Paul Strohmeier. 2013. Magnetic Implant. <http://fkeel.blogspot.de/2013/01/magnetic-implant-sensing.html>
- [31] Paul Strohmeier, Jesse Burstyn, and Roel Vertegaal. 2015. Effects of Display Sizes on a Scrolling Task using a Cylindrical Smartwatch. *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct - MobileHCI '15* (2015), 846–853. <https://doi.org/10.1145/2786567.2793710>
- [32] Paul Strohmeier and Kasper Hornbæk. 2017. Generating Haptic Textures with a Vibrotactile Actuator. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*. ACM Press, New York, New York, USA, 4994–5005. <https://doi.org/10.1145/3025453.3025812>
- [33] Yuriko Suzuki and Minoru Kobayashi. 2005. Air jet driven force feedback in virtual reality. *IEEE computer graphics and applications* 25, 1 (2005), 44–47.
- [34] Radu-Daniel Vatavu, Annette Mossel, and Christian Schönauer. 2016. Digital Vibrons: Understanding Users' Perceptions of Interacting with Invisible, Zero-weight Matter. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '16)*. ACM, New York, NY, USA, 217–226. <https://doi.org/10.1145/2935334.2935364>
- [35] Malte Weiss, Chat Wacharamanotham, Simon Voelker, and Jan Borchers. 2011. FingerFlux: near-surface haptic feedback on tabletops. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*. ACM, 615–620.
- [36] Graham Wilson, Thomas Carter, Sriram Subramanian, and Stephen A. Brewster. 2014. Perception of Ultrasonic Haptic Feedback on the Hand: Localisation and Apparent Motion. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 1133–1142. <https://doi.org/10.1145/2556288.2557033>
- [37] Junhan Zhou, Yang Zhang, Gierad Laput, and Chris Harrison. 2016. AuraSense: enabling expressive around-smartwatch interactions with electric field sensing. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, 81–86.