

# Headio: Zero-Configured Heading Acquisition for Indoor Mobile Devices Through Multimodal Context Sensing

<sup>1</sup>Zheng Sun, <sup>1</sup>Shijia Pan, <sup>2</sup>Yu-Chi Su, and <sup>1</sup>Pei Zhang

<sup>1</sup>Electrical and Computer Engineering, Carnegie Mellon University

<sup>2</sup>Graduate Institute of Electronics Engineering, National Taiwan University

{zhengs, peizhang}@cmu.edu, shijia.pan@west.cmu.edu, steffi@video.ee.ntu.edu.tw

## ABSTRACT

Heading information becomes widely used in ubiquitous computing applications for mobile devices. Digital magnetometers, also known as geomagnetic field sensors, provide absolute device headings relative to the earth's magnetic north. However, magnetometer readings are prone to significant errors in indoor environments due to the existence of magnetic interferences, such as from printers, walls, or metallic shelves. These errors adversely affect the performance and quality of user experience of the applications requiring device headings. In this paper, we propose Headio, a novel approach to provide reliable device headings in indoor environments. Headio achieves this by aggregating ceiling images of an indoor environment, and by using computer vision-based pattern detection techniques to provide directional references. To achieve zero-configured and energy-efficient heading sensing, Headio also utilizes multimodal sensing techniques to dynamically schedule sensing tasks. To fully evaluate the system, we implemented Headio on both Android and iOS mobile platforms, and performed comprehensive experiments in both small-scale controlled and large-scale public indoor environments. Evaluation results show that Headio constantly provides accurate heading detection performance in diverse situations, achieving better than 1° average heading accuracy, up to 33X improvement over existing techniques.

## ACM Classification Keywords

C.3 Special-purpose and application-based systems: Signal processing systems

## Author Keywords

Orientation; heading; digital compass; ceiling pictures; indoor locationing; indoor navigation; mobile sensing; perspective transformation; task scheduling; geolocation

## INTRODUCTION

As the development of ubiquitous computing technologies, heading information of mobile devices becomes important

and are necessary in numerous ubiquitous applications, ranging from indoor localization and navigation [4, 12, 13, 21], activity recognition [5], augmented reality [11, 23], photographing [18], to mobile gaming [15].

Digital magnetometers, also known as geomagnetic field sensors, are commonly used in mobile devices to provide absolute heading information relative to the earth's magnetic north. However, when used in typical indoor environments, magnetometers always suffer from strong magnetic interferences, such as reinforced concrete structures of the buildings [6], or various indoor metallic objects, such as electronic devices, water pipes, electrical conduits inside walls, or metallic supports of furniture. These magnetic interferences lead to significant heading errors. Previous work has shown that median errors in indoor environments can be as large as 17°, which is 15 times greater than outdoors [22]. In some extreme cases, errors over 40° are also observable [9, 13, 22].

To compensate for heading errors in indoor environments, previous work has explored different approaches. One approach is to integrate magnetometer readings with other sensor readings, especially those of gyroscopes and accelerometers, through sensor fusion [14]. However, since a gyroscope only measures *relative* angular changes, it cannot be used to provide absolute device headings. Another approach attempts to reduce magnetic interferences by using averaging or filtering processes on consecutive magnetometer readings [24]. However, these operations are effective only if the magnetic interferences are temporary. If constant interferences exist, such as when a user sitting in a cubicle area, the magnetometer errors cannot be corrected.

Recent work has also used cameras to evaluate photos with pre-tagged orientation and location information [22]. However, the system simplifies its use cases by making camera perspective requirements when capturing ceiling photos. In addition, it assumes a tagging phase through crowdsourcing would be performed beforehand to estimate the ceiling pattern orientation of *every* building, which can be time-consuming and error-prone. Moreover, these methods require a dedicated back-end server for device heading computations. These design choices make the system difficult to operate and maintain in reality, and significantly limit their accuracy and scalability.

In this paper, we propose Headio, a zero-configured heading sensing system for indoor mobile devices. Headio takes advantage of the observation that, as the development of sus-

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](mailto:permissions@acm.org).  
*UbiComp '13*, September 8–12, 2013, Zurich, Switzerland.  
Copyright © 2013 ACM 978-1-4503-1770-2/13/09...\$15.00.

pended acoustical ceiling systems in modern building constructions, various ceiling objects, such as beams, panel grids, tube lamps, and ventilation ducts, are generally mounted in such a way that their straight edges are either parallel or perpendicular to the orientation of the building, to retain aesthetic neatness and to facilitate construction [22]. Headio detects these visual patterns on ceilings, and uses their straight edges to provide directional references. Therefore, when a user wants to determine her phone’s current heading, Headio would capture ceiling images using the front-facing camera of the mobile phone, and automatically compute the heading by integrating the directions of the detected ceiling edges and the device’s own magnetometer readings.

To determine accurate device headings, Headio senses multi-modal contexts with zero-configurations from the user. First, the system uses geolocation sensors on mobile devices to identify locations of users with building-level accuracy. This location information is then used in the detection of building orientations through online map services (e.g. Google Maps or Bing Maps). Second, since users hold mobile devices, such as smartphones, with arbitrary poses, Headio estimates the users’ phone poses using the gravity sensors of the mobile devices, and rectifies perspective distortions to minimize potential heading errors. Finally, to provide energy-efficient heading sensing, Headio collects various ambient contexts, such as phone placements and ambient indoor luminance, to assess the probability of getting effective ceiling images, and dynamically schedules sensing tasks.

To evaluate Headio in real environments, we implemented the system on both Android and iOS mobile platforms, and performed comprehensive experiments in multiple small-scale controlled and large-scale indoor environments. Evaluation results show that Headio constantly achieves high heading detection accuracy in different indoor scenarios, with better than 1° average heading accuracy. This accuracy is up to 33X improvement over current built-in sensor fusion technologies.

The key contributions of this paper are as follows:

- 1) We propose Headio, an absolute heading sensing system for indoor mobile devices using ubiquitous ceiling patterns. Headio can be readily used on commodity mobile phones, without any special hardware or software requirement.
- 2) We develop multiple ambient sensing techniques in Headio to compensate for users’ imperfect phone poses, to deal with heading ambiguities, and to dynamically schedule sensing tasks, with zero configuration from the users.
- 3) We present experimental results collected in multiple small-scale and large-scale indoor environments, giving a quantitative analysis of the effect from indoor metallic surroundings to the performance of magnetometers in mobile devices, and compare it with the performance of Headio.

The following sections provide a detailed description of the Headio system, along with comprehensive evaluations to show its efficacy. Then we highlight the difference between related work and Headio. Finally, we conclude our work and summarize our contributions.

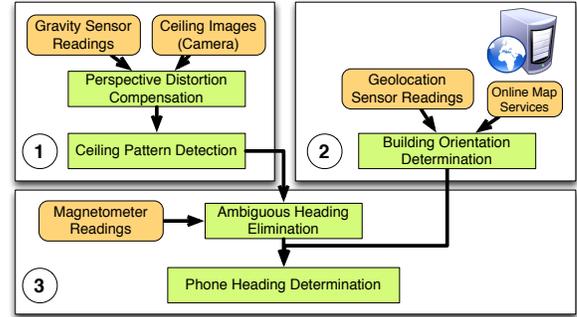


Figure 1. System architecture of Headio. Data from multiple sensory modalities, including gravity sensors, cameras, magnetometers, geolocation sensors, are used in Headio for accurate heading detections.

## SYSTEM OVERVIEW

As shown in Figure 1, the architecture of the Headio system has three main components. First, given the observation that straight edges of ceiling objects are mostly perpendicular or parallel to the building orientation, Headio uses these edges as directional references. In our experiments, we observed that unobstructed ceiling images can be obtained using the front-facing camera when users conduct normal interactions with the device. Headio detects straight edges from these ceiling images using image processing techniques, and computes the orientations of the edges relative to the device.

Headio makes an assumption that most public buildings are rectangular. Though this assumption is not true for some particular buildings, such as the Pentagon, architects A. F. Bemis and M. J. T. Kruger have found out that 90% of modern buildings are predominantly rectangular [19]. Given this assumption, in the second component, Headio uses the geolocation sensor on mobile devices to determine coarse-grained user locations, and uses online map services to determine accurate building orientations. In cases where buildings are not rectangular, the mobile device turns back to its default indoor heading sensing application, such as the digital compass. Next, the building orientations are used to calculate the absolute orientation of the device relative to the earth’s true north. Finally, Headio integrates the orientations of the ceiling edges and that of the building to detect the device headings.

To acquire accurate heading information with zero configurations from the users, Headio addresses a few major technical challenges as follows:

### 1. Obtaining Ceiling Edges From Arbitrary Phone Poses.

To accurately determine device headings, Headio computes orientations of the detected ceiling edges relative to the device. However, due to users’ arbitrary phone placements, ceiling images captured using the device’s front-facing camera may suffer from horizontal heading deviations and perspective distortions, which cause significant heading errors. Headio addresses these issues by using the device’s gravity sensor to estimate the phone’s arbitrary pose, and to compensate for possible perspective distortions and heading deviations.

**2. Getting Absolute Headings From Ceiling Images.** Headio uses ceiling edges as directional references. However,

since both perpendicular and parallel ceiling edges may *concurrently* be present on ceilings, Headio cannot determine the correct device heading from up to four 90°-rotational-symmetric heading ambiguities. To solve this problem, Headio uses the device’s own magnetometer readings as references, and correctly eliminates ambiguities.

**3. Achieving Energy-Efficient Heading Detections.** Headio detects ceiling edges using the front-facing camera of mobile devices. As compared to using magnetometers, the system has the potential risk of consuming significantly more power. In addition, varying distances to lighting sources also results in brightness changes in the captured ceiling image, which make heading detections challenging. To address these challenges, Headio leverages gravity sensors to estimate phone poses and analyzes ambient indoor luminance conditions, to dynamically schedule sensing and computation tasks, without any involvement from the users.

The following sections describe these techniques in details.

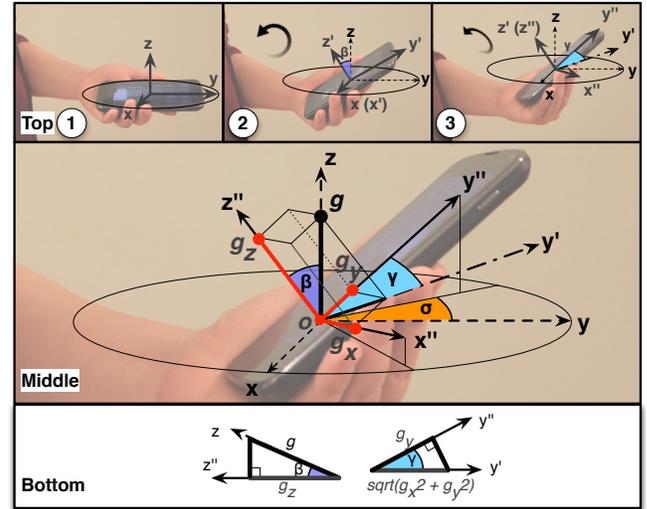
**OBTAINING CEILING EDGES FROM ARBITRARY PHONE POSES**

Headio provides absolute directional references by capturing straight edges on ceilings. However, due to users’ arbitrary phone poses, a minor perspective change of the device’s front-facing camera would yield significant errors in heading detections. To accurately determining device headings, Headio eliminate perspective distortions using two steps. The first step establishes a 3D coordinate transform that maps ceiling objects captured by any arbitrary phone pose to the *ideal* strictly horizontal phone pose. Then, the second step leverages the 3D coordinate transform to correctly rectify a perspective-distorted ceiling image, and detects the orientations of ceiling edges relative to the device.

**Estimating Horizontal Heading Deviations**

Headio leverages the gravity sensors commonly equipped on mobile devices to estimate the 3D pose of the device and horizontal heading deviations. Principles of classical mechanics indicates that every arbitrary 3D pose of a rigid body, i.e. the mobile device in our case, can be decomposed into a sequence of consecutive spatial rotations that starts from an initial reference pose and are parameterized by three Euler angles [8]. To compute the device heading, we define the reference pose as a strictly horizontal *ideal* pose, with an unknown heading on the horizontal plane. This horizontal heading is what Headio will determine eventually, and we name this horizontal pose as the Initial Pose. Let  $X$ ,  $Y$ , and  $Z$  be the three orthogonal axis that form the device’s coordinate system, then the device heading can be uniquely represented by the device’s  $Y$ -axis in Figure 2 (top-1).

The first rotation takes place by rotating the device about its own  $Z$ -axis (vertical) by an angle  $\alpha$  (the first Euler angle), resulting in a horizontal heading change. However, since Headio is meant to determine the device’s horizontal heading, it would always simplify the derivation if we define the Initial Pose in such a way that  $\alpha$  is equal to 0. As shown in Figure 2 (top-2), the second rotation takes place along the device’s own  $X$ -axis (horizontally pointing outwards from the image)



**Figure 2. Illustration of user’s arbitrary phone poses, and their relationships with the Euler angles.** Top row: 1) A strictly horizontal phone pose with the first Euler angle  $\alpha$  (not drawn) being 0. 2) After a rotation about the  $X$ -axis, the second Euler angle  $\beta$  is generated. 3) After another rotation about the  $Z$ - (now  $Z'$ -) axis of the phone coordinate system, the third Euler angle  $\gamma$  is generated. Middle: Using the gravity’s three projections  $g_x$ ,  $g_y$  and  $g_z$ , Headio estimates the phone’s pose by deriving the Euler angles  $\beta$  and  $\gamma$ , as well as the horizontal heading deviation  $\sigma$ . Bottom: Geometric illustrations of the computation of  $\beta$  and  $\gamma$ . The illustration of  $\sigma$  is omitted due to space limit.

by an angle  $\beta$  (the second Euler angle), causing a “pitch” change. Notice that after this rotation, the  $Z$ -axis in the device’s coordinate system is rotated to  $Z'$ -axis, and  $Y$ -axis to  $Y'$ -axis. The last rotation takes places along the device’s new  $Z$ -axis (i.e.  $Z'$ -axis), with an angle  $\gamma$  (the third Euler angle), which reaches the device’s ultimate pose. After the final rotation, the device’s  $Y'$ -axis is rotated to  $Y''$ -axis, as shown in Figure 2 (top-3). Note that the aforementioned rotation sequence following  $Z$ -,  $X$ -, and (again)  $Z$ -axis, is just one possible way of decomposing the pose of a device, yet the described technique can be applied to any rotation sequences. With the three Euler angles  $\alpha$  (equal to 0 in our case),  $\beta$ , and  $\gamma$ , any 3D phone poses can be uniquely represented.

Headio defines the device heading as “*the heading of the projection of the device’s  $Y''$ -axis on the horizontal plane*”. Since with the second rotation, the  $Z$ -axis in the device’s coordinate system is rotated to  $Z'$ -axis (as shown in Figure 2 (top-2)), no longer perpendicular to the horizontal plane in the Initial Pose, the subsequent rotation with the third Euler angle  $\gamma$  would cause the projection of the device’s  $Y''$ -axis to no longer overlap the original  $Y$ -axis, yielding a horizontal heading deviation of the device, denoted as the angle  $\sigma$  in Figure 2 (middle). The heading deviation becomes significant as  $\gamma$  increases, as shown in Figure 3, especially when a camera pose is just slightly apart from the horizontal pose. Figure 4 shows a few examples of common user poses and their corresponding heading deviations. In our experiments, horizontal heading deviations around 10°-30° are commonly seen, which lead to significant heading errors.

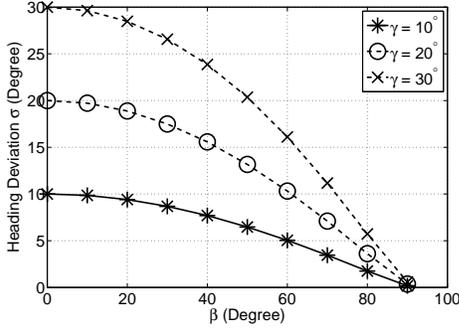


Figure 3. Variation of the heading deviation  $\sigma$  on the horizontal plane as expressed as a function of the second and the third Euler angles  $\beta$  and  $\gamma$ . Note that in our case, the first Euler angle  $\alpha$  equals 0.



Figure 4. Examples of Euler angles  $\beta$  and  $\gamma$  in users' common phone poses, and their resultant heading deviation  $\sigma$ .

### Compensating for Perspective Distortions

To accurately compute the device heading and to compensate for heading deviations, Headio derives  $\beta$ ,  $\gamma$ , and  $\sigma$  using the gravity sensor readings of the mobile device. The gravity sensor readings  $g_x$ ,  $g_y$ , and  $g_z$  are normalized projections of the earth's gravity  $g$  on the device's axes [1]. For example, when the device is horizontally placed, the gravity sensor readings are  $[0, 0, 1]$ . As shown in Figure 2 (bottom), through geometric analysis, the angles can be computed as  $\beta = \text{acos}(\text{Abs}(g_z))$ ,  $\gamma = \text{acos}(\frac{\text{Abs}(g_y)}{\sqrt{g_x^2 + g_y^2}})$ , and

$$\sigma = \frac{\cos(\gamma)}{\sqrt{1-g_x^2}}.$$

Suppose a ceiling object in the world coordinate system is  $[x, y, z]^T$ , according to principles of homogeneous coordinate transforms [10], its corresponding coordinates in the phone's coordinate system ( $P$ ) would be

$$[x_P, y_P, z_P, 1]^T = M(\alpha, \beta, \gamma) \cdot [x, y, z, 1]^T, \quad (1)$$

where  $M(\alpha, \beta, \gamma)$  is the 3D rotation matrix uniquely defined by the combination of the three Euler angles. Therefore, the ceiling object would be projected onto the camera's pixel coordinate system through the following projective transform [10]

$$[u_P, v_P, 1]^T \sim \begin{bmatrix} \frac{f}{s_x} & 0 & o_x & 1 \\ 0 & \frac{f}{s_y} & o_y & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} [x_P, y_P, z_P, 1]^T, \quad (2)$$

where  $u_P$  and  $v_P$  are coordinates of the objects in the camera's pixel coordinate system,  $f$  the focal length of the camera,  $s_x$  and  $s_y$  the scale factors that map camera's sensor dimensions to pixel dimensions,  $o_x$  and  $o_y$  the image center

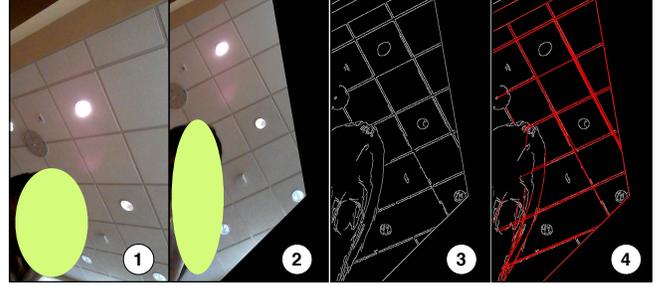


Figure 5. 1) A perspective distorted ceiling image. 2) After perspective transformation, perpendicular and parallel ceiling grids are restored. 3) Using the Canny algorithm to detect ceiling edges. 4) Using the Hough transform to detect straight lines. Nota that, the user's head was captured in the image, which caused noise in line detections. All images are screenshots directly taken by Headio running on a Nexus 4 smartphone, except that for the double-blinded review, we purposely marked out the user's faces on the first two colorful images.

coordinates of the camera, and “ $\sim$ ” the homogenous transformation operation [10].

Equation (1) and (2) defines a direct mapping relationship between a ceiling object  $[x, y, z]^T$  in the world coordinate system to its correspondence  $[u_P, v_P]^T$  on an image that is taken by the device with an arbitrary pose. This relationship is completely defined by the three Euler angles, which are derived using the gravity sensor readings. Suppose there is an *imaginary* ceiling object, its image can be derived using the Euler angles (note the first Euler angle  $\alpha = 0$  in our case) as

$$[u_P, v_P, 1]^T \sim A \cdot M(0, \beta, \gamma) \cdot [x, y, z, 1]^T, \quad (3)$$

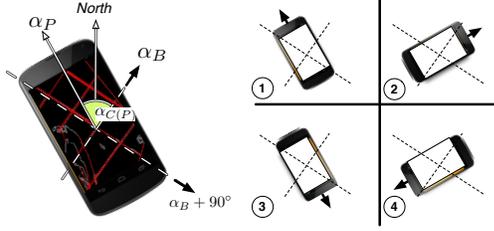
where  $A$  denotes the 3-by-4 matrix in Equation (2). Similarly, to compensate for the heading deviation, the *same* ceiling object can be mapped to another image, which is taken when the device is horizontally placed but with a horizontal heading  $\sigma$ . In this case, the first Euler angle  $\alpha$  is equal to the heading deviation  $\sigma$ , and both  $\beta$  and  $\gamma$  are equal to 0. This mapping relation can be expressed as

$$[u_D, v_D, 1]^T \sim A \cdot M(\sigma, 0, 0) \cdot [x, y, z, 1]^T, \quad (4)$$

where  $u_D$  and  $v_D$  are coordinates of the *same* ceiling objects on the image taken when the device is horizontally placed but with a horizontal heading  $\sigma$ .

Considering that all the unknowns in matrix  $A$  are intrinsic camera parameters that can be determined using the camera's specifications,  $[x, y, z, 1]^T$  can be cancelled out in the two equations, and  $\beta$ ,  $\gamma$ ,  $\sigma$  are derived using gravity sensor readings, Equation (3) and (4) define a readily solvable perspective transformation operation between a ceiling image taken with an arbitrary device pose to its rectified version without perspective distortions. Figure 5 (1) and (2) show the effect of the perspective rectification operation.

After rectifying the perspective distortions in the ceiling image, Headio detects visible edges by processing each image frame using the Canny edge detection algorithm, followed by a probabilistic Hough transform [10], as shown in Figure 5 (3) and (4).



**Figure 6.** Left: Relations between the orientation of ceiling edges  $\alpha_B$  and that of the mobile device  $\alpha_P$ , both relative to the earth's true north. Right: The four heading ambiguities caused by the existence of  $90^\circ$ -rotational-symmetric ceiling edges.

## GETTING ABSOLUTE HEADINGS FROM CEILING IMAGES

Ceiling images usually contain diverse and complex objects, such as beams, grids, round lamps, or ceiling fans. These objects may feature curvy edges, or straight edges that have random angles. To correctly determine the building orientation relative to the mobile device, Headio analyzes dominant angles of the detected straight edges. This is achieved through a two-round linear search. In the first round, if more than  $p\%$  detected lines have a similar angle with  $\pm b^\circ$  tolerance, this angle is considered as a candidate dominant angle. Then, a second-round search is conducted among the existing candidate angles, looking for any pairs of angles that are  $90^\circ$  apart. If the search succeeds, Headio considers the image contains perpendicular lines, and uses the smaller angle in the pair as the dominant angle; otherwise, the system considers that only parallel lines exist in the image, and uses the angle in the first round that had the most lines reside on as the dominant angle. In the experiments, we empirically set  $p$  and  $b$  as 20% and  $2^\circ$ , respectively, and found that this approach achieved stably good performance in detecting dominant angles.

### Determining Absolute Building Orientations

To acquire building orientations, Headio uses the geolocation sensor on mobile devices, to make a building-level estimate of the current user location. The map view that contains the building image is then obtained through the online map services, such as the Google Map, and processed using the Canny and Hough transform to estimate the building orientations relative to the earth's true north. Due to the directional difference between the earth's *magnetic* north and *true* north, we also compensate for this difference by taking into account the local magnetic declination [2].

Let  $\alpha_B$  denote the orientation of the building, and  $\alpha_{C(P)}$  the dominant orientation of the ceiling edges relative to the smartphone, then the device heading  $\alpha_P$  relative to the earth's true north would be

$$\alpha_P = \text{Mod}(\alpha_B - \alpha_{C(P)}, 360^\circ), \quad (5)$$

where  $\text{Mod}(\cdot)$  represents the computation of remainders after division. As shown on Figure 6, due to the existence of *both* parallel and perpendicular lines on the ceiling, the system cannot determine the correct device heading from up to four  $90^\circ$ -rotational-symmetric heading ambiguities. For example, if the dominant direction of the ceiling edges relative

to the phone is  $50^\circ$ , and the building orientation is  $68^\circ\text{E}$ , Equation 5 would indicate that the device heading is  $18^\circ\text{N}$  (i.e.  $68^\circ\text{E} - 50^\circ$ ), whereas the actual device heading could be  $18^\circ\text{N}$ ,  $108^\circ\text{E}$ ,  $198^\circ\text{S}$ , or  $288^\circ\text{W}$ .

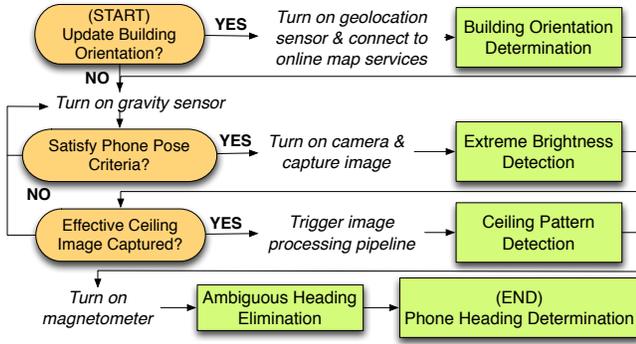
### Handling Directional Ambiguities

To eliminate the heading ambiguities, Headio uses the magnetometer readings as references, and selects the heading ambiguity that is closest to the magnetometer reading as the final device heading. For example, in the case above, if the magnetometer reading is  $36^\circ\text{NE}$ , Headio would choose  $18^\circ\text{N}$  as the final device heading. It should be noted that, due to magnetic interferences, magnetometer readings can be erroneous. This means if the magnetometer error is greater than  $\pm 45^\circ$ , Headio will choose an incorrect heading value and generate a significant heading error. However, during our experiments in both small-scale and large-scale experiments, we found that Headio was able to achieve accurate heading detections in the majority (i.e. over 95%) of common indoor conditions.

### ACHIEVING ENERGY-EFFICIENT HEADING DETECTION

Headio detects device headings through the use of front-facing camera. As compared to traditional magnetometer sensors, the downside is the potential increase of power consumptions. To address this problem, we have designed and implemented a dynamic sensing task scheduling strategy to improve the energy efficiency of Headio, which schedules sensing and computation tasks based on ambient contexts. Since ceiling objects can be hardly captured with a huge perspective distortion, Headio first uses the gravity sensor readings to estimate the device pose, and the scheduling process proceeds only if the current phone pose satisfies the pose criteria. As will be shown in our evaluations, we found that when the Euler angle  $\beta$  and  $\gamma$  exceed  $20^\circ$  and  $30^\circ$ , respectively, the camera held by a user with a normal body height will usually capture objects on walls or nearby objects, such as shelves along aisles in a pharmacy. Therefore, in our experiments, we empirically set the pose criteria to be: 1)  $\text{Abs}(\beta) \leq 20^\circ$ , and 2)  $\text{Abs}(\gamma) \leq 30^\circ$ .

If the pose criteria is satisfied, the device's camera will be turned on, and image frames will be captured and processed continuously. In our experiments, we noticed that a strong luminance condition, such as when a tube lamp with strong light is captured in the image, or a weak luminance condition, could dramatically increase the brightness of ceiling images, causing the ceiling patterns to be buried in dark backgrounds and become completely invisible. In such extreme cases, due to the lose of colorful details, running the computationally expensive Canny and Hough transforms will not detect correct ceiling edges. Therefore, Headio first conducts a relatively simple and fast color analysis of each image, and eliminates images with extremely high or low brightness. This is achieved by converting the image from the RGB space to the HSB (i.e. Hue, Saturation, and Brightness) space, with the last dimension representing the normalized brightness, and then analyzing the brightness distribution of the image. If the numbers of pixels that are over-bright or over-dark exceed  $d\%$ , the image is considered to have extreme brightness, and the scheduling process stops. Otherwise, the scheduler would



**Figure 7.** The dynamic sensing task scheduling strategy of Headio. Note that sensing and computation tasks with heavier power consumptions are always scheduled behind lighter ones, so that unnecessary power consumptions can be avoided.

enable the core Canny and Hough transforms, and detect device headings using the ceiling image. In our experiments, we set  $d$  and the threshold of determining over-brightness and over-darkness to be 5%, 0.02, and 0.98, respectively. Using this scheduling strategy, we found that most of the unnecessary but power-hungry sensing and computation tasks can be avoided, and more experimental results are discussed in the evaluation section.

### SYSTEM IMPLEMENTATION

To evaluate the system, we implemented Headio on both the Android and the iOS mobile platforms, and tested it on various devices, including Nexus 7 (tablet), Nexus 4, Galaxy Nexus, iPhone 4 and 4S. To compute the matrix  $A$  discussed in the perspective transformation section, we collected the related information, including the focal length of the front-facing camera and the pixel dimensions of the image, from the EXIF metadata of the pictures. To get accurate building orientations relative to the earth’s true north, instead of the magnetic north, we determine the magnetic declination of our local area using the online public database provided by the National Geophysical Data Center [2], and compensate for the magnetic declination accordingly. The perspective transformation and all the image rendering tasks were implemented using OpenCV [3].

Headio runs the aforementioned sensing tasks scheduler in the background, and processes each incoming image frame continuously. When the criteria of heading detections are satisfied, device headings are determined opportunistically and automatically, and shown on the screen of the device. The entire Headio project contains a total of 2,700+ lines of Java code, and 2,000+ lines of Objective-C code.

### EVALUATIONS

Headio determines device headings using ubiquitous ceiling patterns. To understand the boundary of system performance, we perform three types of experiments: 1) we conduct experiments to evaluate the rectification of perspective distorted ceiling images; 2) we provide end-to-end performance analysis of the system in multiple small-scale and large-scale experiments; 3) we discuss the energy efficiency of using the



**Figure 8.** Experimental setup: 1) The movable plastic platform we built to facilitate the measurement of ground-truth headings. Two Fluke laser distance meters were used to keep the platform parallel to the building orientation. 2) Small-scale experiment location: a student activity room. 3) and 4) are the locations of the two large-scale experiments: an office building, and a CVS pharmacy.

sensing task scheduling technique. We also compare power consumptions of Headio with other common mobile applications.

To keep the actual phone heading unchanged and facilitate the measurement of ground truths, we built a movable plastic platform, with two Fluke laser distance meters mounted in parallel, as shown in Figure 8. During our experiments, we placed the device in parallel to the platform’s orientation, and by keeping the distance reading of the two laser meters the same, we guaranteed the platform to be parallel to the orientation of the building.

### Performance of Handling Arbitrary Phone Poses

In our tests with participants, we found that, when indicated to take ceiling images, average users would hold the phone with limited Euler angle ranges. To be specific, both  $\beta$  and  $\gamma$  are commonly seen within the range from  $0^\circ$  to  $30^\circ$ . To evaluate the effect of users’ arbitrary phone poses to device headings, we perform experiments in our student activity room, which features visible ceiling grids. Before the experiment, we measured the ground-truth building orientation at several outdoor locations using a compass. During the experiment, we measured the building’s orientation by accessing the Google Map. To measure the ground truths of phone poses and to prevent magnetic interferences from metal, we fixed the phones using a plastic vise.

The system performance was tested under different Euler angle combinations. As shown in Figure 9, with all different  $\gamma$  values, Headio achieved about  $0.5^\circ$  heading accuracy. However, when the value of  $\beta$  increases, the heading errors increase significantly. This is due to the fact, as shown in Figure 2,  $\beta$  represents the “pitch” change of the phone pose, and when  $\beta$  is great, the camera becomes more likely to capture indoor objects that are not on ceilings, which decrease the accuracy of detecting the orientation of ceiling objects. To achieve constantly acceptable performance, we use the sensing tasks scheduling technique to tighten the pose criteria so

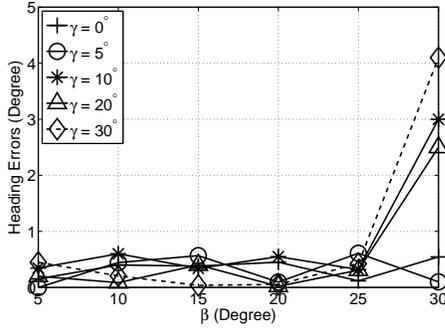


Figure 9. Performance of handling arbitrary phone poses. Experimental results show that with all different  $\gamma$  values, Headio achieved about  $0.5^\circ$  heading accuracy. However, as  $\beta$  increases, the heading errors increase significantly, due to the unintentional capture of indoor objects that are not on ceilings.

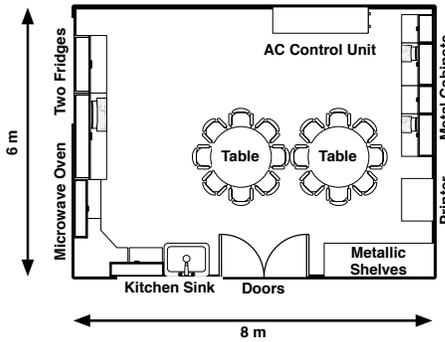


Figure 10. Floor plan of the room where we did the small-scale controlled experiment. The room features various common indoor metallic objects.

that Headio is disabled when  $\beta$  is bigger than  $20^\circ$  in our end-to-end experiments.

### End-to-End System Performance

To evaluate the performance boundary of Headio, we conducted end-to-end system evaluation in various indoor environments. Figure 8 shows the three locations where we did the experiments.

#### Small-Scale Finer-Grained Experiment

To understand how common indoor objects affect the performance of Headio as well as the magnetometer sensors in off-the-shelf mobile devices, we first performed experiments in a student activity room, which featured rich magnetic interferences, such as fridges, microwave ovens, and metallic shelves, as shown in Figure 10. We evenly divided this 6m-by-8m room into a 8-by-20 grids, with roughly 0.75m (width) and 0.40m (length) spatial granulations. At each grid point, we tested a Nexus 4 mobile phone on the movable platform to measure the device headings, and compare the performance of Headio with that of the magnetometer reading with the use of Android’s built-in sensor fusion techniques. Similar to the previous experiments, ground-truth building orientations are measured outside of the building using a compass, and during the experiments, the buildings orientation was obtained through the Google Map.

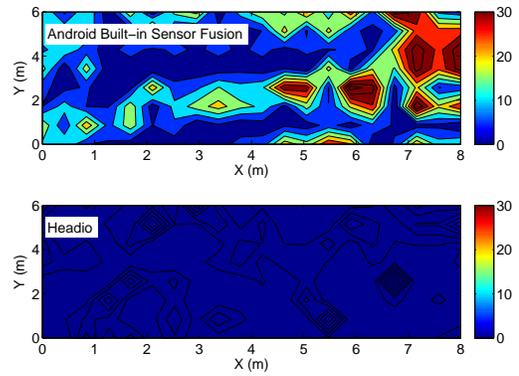


Figure 11. Distribution of heading errors (unit: degree). Magnetic interferences in the room (shown on Figure 10), i.e. the AC control unit, the tables, and metal cabinets, have caused significant negative effects to the magnetometer readings, whereas Headio’s performance is robust against such strong interferences.

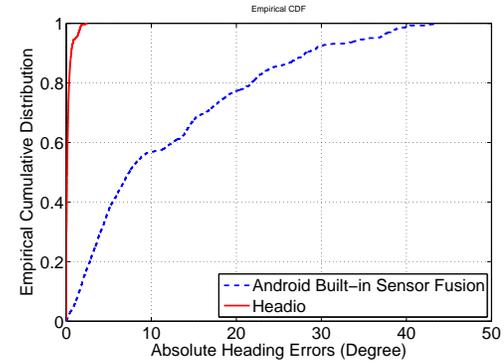
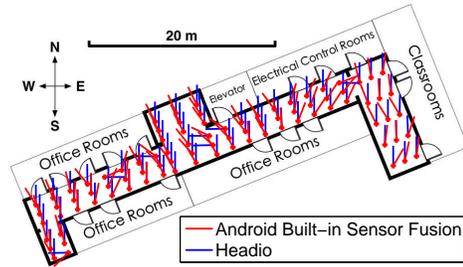


Figure 12. Empirical cumulative distribution of heading errors in the small-scale experiment. Quantitatively, Headio achieved  $0.2^\circ$  (median),  $0.3^\circ$  (75th-percentile), and  $1.4^\circ$  (95th-percentile) heading accuracies, as compared to  $7.4^\circ$ ,  $18.7^\circ$ , and  $37.1^\circ$  when using the magnetometer sensor with the latest Android’s built-in sensor fusion technique. These results are 26X to 37X improvements.

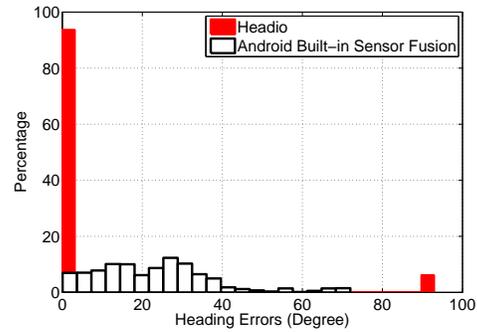
Figure 11 shows the contour distribution of the heading errors. From the figures, it is clear to see a match between the distribution of the error to that of the magnetic interferences in the room. Specifically, the interferences including the AC control unit, the metallic supports of the tables, and the metal cabinets, have caused significant negative effects to the magnetometer readings, whereas Headio’s performance is robust against such strong interferences. To quantitatively analyze the performance improvements of Headio, we plotted the empirical cumulative distribution of heading errors, as shown in Figure 12. Experimental results indicates that Headio achieved  $0.2^\circ$  (median),  $0.3^\circ$  (75th-percentile), and  $1.4^\circ$  (95th-percentile) heading accuracies, as compared to  $7.4^\circ$ ,  $18.7^\circ$ , and  $37.1^\circ$  when using the magnetometer sensor with the latest Android’s built-in sensor fusion technique. These results are 26X to 37X improvements.

#### Large-Scale Coarse-Grained Experiments

To evaluate the performance of Headio in more open and diverse indoor scenarios, we performed two large-scale experiments.



(a) Heading errors at each grid point along the hallway



(b) Distribution of heading errors

**Figure 13. Performance of device heading detection in the large-scale experiment (hallway). Note in Figure (b) that though Headio achieved accurate performance in most cases, there are occasionally large errors in Headio caused by incorrect eliminations of heading ambiguities.**

**The Hallway Experiment.** The first experiment was performed along a hallway of one office building in our university, which had a  $68^\circ$  true building orientation. To measure ground-truth directions the movable platform shown in Figure 8 was used, and the laser distance meters were used to guarantee orientational correctness. One Nexus 4 phone was placed on the platform and pointed towards the true north of the earth. Figure 13(a) shows the heading error distribution throughout the hallway, as well as the quantitative statistics of the error distribution. As shown in the figure, the magnetometer readings suffered from significant magnetic interferences. To be specific, at locations in the West wing of the building, magnetometer readings experienced a westward directional bias, whereas in the East wing of the building, especially at locations outside of the electrical control rooms, a significant eastward bias was commonly observable. These results indicate that the distribution of magnetic errors are location-dependent and not uniformly distributed, thus an automatic direction tagging operation cannot perform accurately.

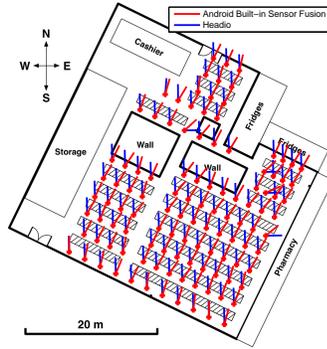
As a comparison, in most case Headio was able to accurately detect device headings, and pointed upwards to the correct direction (i.e. the direction of the true north on the map). However, we also observed there were occasionally large errors in Headio caused by incorrect eliminations of heading ambiguities. As shown in Figure 13(b), these large errors are  $90^\circ$  apart from the correct direction, and were found at locations close to the elevator room. However, such large errors in Headio only occurred at 6% of the grid points in our experiments, and in the majority of the locations (i.e. 94%), Headio performed stably and achieved  $1.1^\circ$  median heading accuracy. As compared to the  $27.1^\circ$  median accuracy of using the Android’s built-in sensor fusion technique, this is equivalent to a 25X improvement.

**The Pharmacy Experiment.** To evaluate the performance of Headio enabling the corrections of perspective distortions, we performed the second experiment in a local CVS pharmacy, which had a  $26^\circ$  true building orientation. To largely reduce intrusiveness to the normal business of the pharmacy, the movable platform was not utilized. Instead, one student walked along aisles of the pharmacy, following straight paths,

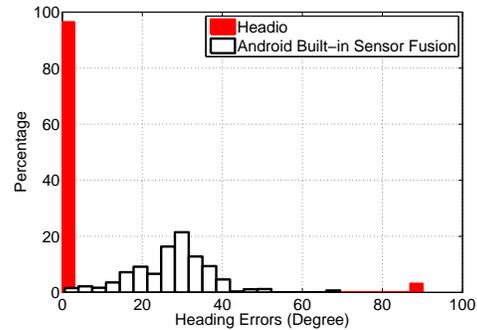
and used Headio to determine device headings. During the experiment, the student held the phone with a roughly horizontal placement as long as ceiling patterns could be successfully captured in the viewfinder of the camera. The student were also indicated that a strict upward-facing perspective was not necessary, since Headio would be able to correct it. At each step, a device heading detection was made in Headio, and logged in the device. After the experiment, we matched the results to locations given the student’s stride length and the aisle IDs. When a person is walking, the accelerometer picks up signals from the person movement. In the current implementation of Headio, we use a motion detection approach to detect a user’s movement, and provide the heading service when the user’s movement is within a certain range.

Figure 14(a) shows the distribution of heading errors of Headio (blue). The performance of using the magnetometer sensor with Android’s built-in sensor fusion technique is also shown (red). As compared to the first experiment along the hallway, we observed significantly stronger magnetic interferences in the pharmacy, especially in the areas of pharmacy shelves. Due to the metallic materials of the shelves, magnetometer readings were predominantly biased towards the shelves, which was about  $29.9^\circ$  eastwards. These interferences caused the magnetometers to perform poorly, with a  $28.5^\circ$  median heading error. As a comparison, since Headio leveraged ceiling patterns to obtain accurate directional references, it was robust against ambient magnetic interferences. Headio’s median heading error was  $0.85^\circ$ , which is equivalent to a 33X improvement.

Figure 14(b) shows the statistics of the heading errors. While the hallway experiment had a flatter but broader error distributions when using the Android’s built-in sensor fusion technique, the pharmacy experiment had a more biased and skewed error distribution. As discussed earlier, this is primarily due to the strong interferences from the metallic shelves. Similar to the results of the hallway experiment, because of the incorrect elimination of heading ambiguities, we observed occasional large errors using Headio in 3% out of all tested locations inside the pharmacy. However, we found these sig-



(a) Heading errors at each grid point in the local pharmacy



(b) Distribution of heading errors

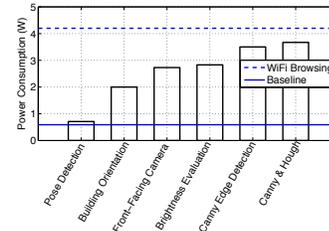
**Figure 14. Performance of device heading detection in the large-scale experiment (pharmacy). Note in Figure (b) Headio’s occasional large errors caused by incorrect eliminations of heading ambiguities. Also note that, as compared to errors in the hallway experiment, the error distribution in the pharmacy is not uniformly distributed, due to the widely predominant interferences from the metallic shelves.**

nificant errors occurred only in rare cases, such as the food sections where refrigerators were placed, and in most of locations (i.e. 97%), Headio constantly provided accurate heading detections.

### Performance of Energy Consumption

Headio detects ubiquitous ceiling patterns using the front-facing cameras of mobile devices. To improve the energy efficiency of Headio, we designed and implemented a sensing task scheduler that dynamically enables or disables sensing and computation tasks based on ambient contexts. In this section, we evaluate the performance of Headio in terms of power consumption. We performed the experiments using a Galaxy Nexus smartphone, and measured the power consumptions using an oscilloscope. For each test, we measured the power consumption by running Headio continuously for 10min, and we compute the average power consumption of each test. The experimental results are shown in Figure 15. As a baseline, we first measure the power consumption of the devices being idle, i.e. dimming the screen and shutting down all the application processes. This yields a baseline power consumption at 0.59W. Then we evaluate the performance of Headio. When the front-facing camera is enabled, Headio consumes 2.73W, whereas the gravity sensors consumes only 0.71W, which is 3X less than enabling the camera. This result justifies Headio’s strategy of using the gravity sensors to estimate phone poses before using the camera, to avoid unnecessary energy consumptions. Our experimental results also show that, in addition to the front-facing camera, the Canny and Hough transforms are also power-hungry. To be specific, when both the transforms are enabled, which are necessary to detect ceiling edges, Headio consumes 3.67W in total, whereas the simple and fast brightness evaluation task consumes only 2.83W, which is 30% more energy-efficient.

As a comparison, we also measured the power consumption of browsing webpages using WiFi, which consumed 4.2W. This result indicate that, even enabling all the sensing and computational tasks (i.e. the gravity sensor, the camera, the Canny and Hough transforms), Headio still consumes 14% less power than WiFi.



**Figure 15. Energy consumptions of technical components in Headio. Using the sensing task scheduling technique, Headio is able to reduce unnecessary power-hungry tasks, such as running the camera and doing the Canny and Hough transforms. Note that, the total power consumptions of the Canny and Hough transforms also include that of enabling the camera.**

### RELATED WORK

Magnetometer sensors have been widely used in ubiquitous applications for finding correct device headings. These applications include indoor locationing [4, 12, 13], activity recognition [5], and augmented reality [23, 11]. However, due to the common existence of indoor magnetic interferences, device headings suffer from significant errors.

To improve the accuracy of heading determination, two major approach have been proposed in previous work. The first approach combines magnetometer readings with readings of additional sensors, especially gyroscopes, through the use of Kalman filters [14]. However, since the gyroscope only measures *relative* angular changes, this approach requires an accurate initial device heading, which may be error-prone due to systematic biases. Moreover, the gyroscopes also drifts from correct directions quickly if subsequent correctness cannot be provided. In reality, this issue may lead to significant heading errors. The second approach attempts to use the time-domain averaging operation to reduce the magnetic interferences. However, this approach is effective only if the magnetic interferences are temporary. If constant interferences exist, such as a student sitting on a chair that has metallic supports. Other similar approaches have attempted to use location changes of users to filter out heading noises. However, as we have shown in the large-scale pharmacy experiments, if

a user wanders in a large pharmacy where the magnetic field is predominantly biased towards metallic shelves, reducing magnetic interferences is difficult. As opposed to using filtering techniques, Headio avoids these problems by using ubiquitous and invariant visual patterns on ceilings. Since ceiling patterns are unrelated to magnetic fields, Headio is not affected by magnetic interferences and able to provide accurate device headings.

Recent work has shown that visual contexts (i.e. floor images) can be used to determine user locations [4]. Similarly, previous work has also used cameras to evaluate ceiling photos with pre-tagged orientation information on mobile phones or robots [7, 22]. However, their work has strict camera perspective requirements, making the systems difficult to use in ubiquitous applications. Moreover, the systems require information tagging (either through crowdsourcing or individual) of every building. These limitations reduce the accuracy of the system, and significantly restrict their scalability. Other projects have applied inertial measurement unites and vision fusions in visual odometry pipelines, or have generated interactive visual tours and 2D floor plans, but both limited by Manhattan world assumptions [16, 17]. Spartacus is a mobile interaction system with directional knowledge [20]. However, since Spartacus deals with only relative orientations of devices, absolute headings of devices are still unknown. As a comparison, Headio determines the relative orientation through intrinsic properties of ceiling patterns and building orientation using online map services that are already commonly used on mobile devices. This removes any additional hardware and deployment requirements. In addition, due to the technique of rectifying the distortions in images, Headio is able to detect device headings without any perspective restrictions.

## CONCLUSION

This paper presents Headio, a zero-configured heading detection system for mobile devices. Unlike traditional use of magnetometer sensors to acquire orientation information, Headio determines device headings by using computer vision techniques to extract intrinsic directional patterns on ceiling in modern public buildings. Since ceiling patterns are unrelated to the magnetic field, Headio is robust against magnetic interferences plaguing current techniques. To reduce user involvements and improve energy efficiency, Headio also utilizes multimodal sensing techniques to dynamically schedule sensing tasks. We provide a comprehensive evaluation of the Headio system in various indoor conditions. Our experimental evaluations show that Headio reduces error to less than 1°, a more than 33X improvement when compared to existing magnetometers with built-in sensor fusion techniques. This new paradigm of accurate heading sensing will enable numerous emerging ubiquitous computing applications.

## ACKNOWLEDGEMENT

This work was partially the results of the generous support from National Science Foundation grant number 1149611.

## REFERENCES

1. Motion Sensors of Android Platform, [http://developer.android.com/guide/topics/sensors/sensors\\_motion.html](http://developer.android.com/guide/topics/sensors/sensors_motion.html). 2013.

2. NOAA's Magnetic declination database, <http://www.ngdc.noaa.gov/geomag-web/#declination>. 2013.
3. OpenCV, <http://opencv.org/>. 2013.
4. M. Azizyan, I. Constandache, and R. Roy Choudhury. SurroundSense: mobile phone localization via ambience fingerprinting. In *Proceedings of the 15th annual international conference on Mobile computing and networking, MobiCom '09*, pages 261–272, New York, NY, USA, 2009. ACM.
5. T. Choudhury, G. Borriello, and S. Consolvo. The Mobile Sensing Platform: An Embedded Activity Recognition System. *IEEE Pervasive*, pages 32–41, 2008.
6. M. Faifer, R. Ottoboni, and S. Toscani. A compensated magnetic probe for steel fiber reinforced concrete monitoring. In *Sensors, 2010 IEEE*, pages 698–703, 2010.
7. T. Fukuda, Y. Yokoyama, F. Arai, K. Shimojima, S. Ito, Y. Abe, K. Tanaka, and Y. Tanaka. Navigation system based on ceiling landmark recognition for autonomous mobile robot -position/orientation control by landmark recognition with plus and minus primitives. In *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, volume 2, pages 1720–1725 vol.2, 1996.
8. H. Goldstein. The Euler Angles. In *Classical Mechanics*, pages 143–148. Addison-Wesley, 1980.
9. D. Gusenbauer, C. Isert, and J. Krosche. Self-Contained Indoor Positioning on Off-The-Shelf Mobile Devices. In *Proceedings of the 2010 International Conference on Indoor Positioning and Indoor Navigation*, number September, pages 1–9. IEEE, 2010.
10. R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2nd edition, 2003.
11. B. Hoff and R. Azuma. Autocalibration of an electronic compass in an outdoor augmented reality system. In *Augmented Reality, 2000. (ISAR 2000). Proceedings. IEEE and ACM International Symposium on*, pages 159–164, 2000.
12. W. Kang, S. Nam, Y. Han, and S. Lee. Improved heading estimation for smartphone-based indoor positioning systems. In *Personal Indoor and Mobile Radio Communications (PIMRC), 2012 IEEE 23rd International Symposium on*, pages 2449–2453, 2012.
13. T. King, S. Kopf, T. Haenselmann, C. Lubberger, and W. Effelsberg. Compass: A Probabilistic Indoor Positioning System Based on 802.11 And Digital Compasses. In *Proceedings of the 1st International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*, number September, pages 34–40. ACM, 2006.
14. S. O. H. Madgwick, A. J. L. Harrison, and R. Vaidyanathan. Estimation of IMU and MARG orientation using a gradient descent algorithm. In *Rehabilitation Robotics (ICORR), 2011 IEEE International Conference on*, pages 1–7, 2011.
15. W. Piekarski and B. Thomas. ARQuake: The Outdoor Augmented Reality Gaming System. *Communications of the ACM*, 45(1):36–38, 2002.
16. A. Sankar and S. Seitz. Capturing indoor scenes with smartphones. In *Proceedings of the 25th annual ACM symposium on User interface software and technology, UIST '12*, pages 403–412, New York, NY, USA, 2012. ACM.
17. O. Saurer, F. Fraundorfer, and M. Pollefeys. Homography based visual odometry with known vertical direction and weak Manhattan world assumption. In *Vicomor Workshop at IROS 2012*, 2012.
18. D. Spinellis. Position-Annotated Photographs: A Geotemporal Web. *IEEE Pervasive Computing*, 2(2):72–79, Apr. 2003.
19. P. Steadman. Why Are Most Buildings Rectangular? *Architectural Research Quarterly*, 10(02):119–130, June 2006.
20. Z. Sun, A. Purohit, R. Bose, and P. Zhang. Spartacus: Spatially-Aware Interaction for Mobile Devices Through Energy-Efficient Audio Sensing. In *The 11th MobiSys*, pages 263–276, 2013.
21. Z. Sun, A. Purohit, K. Chen, S. Pan, T. Pering, and P. Zhang. PANDAA: physical arrangement detection of networked devices through ambient-sound awareness. In *Proc. the 13th ACM Ubicomp, UbiComp '11*, pages 425–434. ACM, 2011.
22. Z. Sun, A. Purohit, S. Pan, F. Mokaya, R. Bose, and P. Zhang. Polaris: getting accurate indoor orientations for mobile devices using ubiquitous visual patterns on ceilings. In *Proc. the 12th ACM HotMobile, HotMobile '12*, pages 14:1—14:6. ACM, 2012.
23. S. You, U. Neumann, and R. Azuma. Orientation tracking for outdoor augmented reality registration. *Computer Graphics and Applications, IEEE*, 19(6):36–42, 1999.
24. M. Youssef, M. Yosef, and M. El-Derini. GAC : Energy-Efficient Hybrid GPS-accelerometer-compass GSM Localization. In *Proceedings of the IEEE Globecom 2010*, 2010.