UNIVERSITY OF CALIFORNIA
SANTA CRUZ

# BENCHMARKING THE ACCURACY OF INERTIAL SENSORS IN CELL PHONES

A thesis submitted in partial satisfaction
of the requirements for the degree of

MASTER OF SCIENCE

in

ELECTRICAL ENGINEERING

by

**Bin An**

March 2012

The Dissertation of Bin An
is approved:

Professor Roberto Manduchi, Chair

Professor Peyman Milanfar

Professor James Davis

Dean Tyrus Miller
Vice Provost and Dean of Graduate Studies

# Table of Contents

# List of Figures

# List of Tables

**Abstract**

BENCHMARKING THE ACCURACY OF

INERTIAL SENSORS IN CELL PHONES

by

Bin An

Many ubiquitous computing applications rely on data from a cell phone's inertial sensors. Unfortunately, the accuracy of this data is often unknown, which impedes predictive analysis of applications that require high sensor accuracy (e.g., dead reckoning). This work focuses on benchmarking the accuracy of the accelerometers and gyroscopes on a cell phone. The cell phones are attached to a robotic arm, which provides ground truth measurements. The misalignment between the cell phone's and the robotic arm's reference systems is computed using Horn's algorithm for closed-form absolute orientation estimation. Two cell phones (Apple's iPhone 4 and Nokia's N97) have been tested, and results are provided in terms of random noise, error bias, error bias correlations, and orientation reconstruction error.

This work is dedicate to my advisor,

Professor Roberto Manduchi,

whose help and support made this project possible.

# Chapter 1

# Introduction

Recent years have witnessed great interest in the use of inertial sensors embedded in cell phones. Nowadays most smartphones contain 3-axis MEMS accelerometers which can be used to estimate the inclination of the phone with respect to the vertical when the phone is static or under linear motion, and 3-axis rate gyroscopes that allow for direct measurement of angular velocities. Data from cell phone inertial sensors can be used for a variety of ubiquitous computing applications. Perhaps the simplest use of an accelerometer is as a pedometer; in this case, the sensor only needs to detect footsteps, corresponding to strong acceleration values. Other applications call for a higher level of accuracy. For example, gesture recognition [15, 2] requires precise motion and inclination measurements. Another example is given by the use of inertial sensors to assist visual tracking [14, 3, 9, 17] and for assisted camera pose estimation [7, 8]. An excellent overview of inertial sensors is provided in [18].

An important component of my research is the development, implementation

and testing of algorithms for camera orientation estimation via inertial sensors, which supplement image-based estimation. It is well known that a 3-axes accelerometer can only provide a partial estimation of the orientation. Let the natural coordinate frame be defined such that its $X$ axis and $Y$ axes are in the horizontal plane and the $X$ axis points to north, while the $Z$ axis points down (in the direction of gravity). The cell phone's reference system is defined such that its $X$ axis is orthogonal to the screen (pointing towards the back of the phone), its $Y$ axis is aligned with the shortest side of the screen (pointing right), and its $Z$ axis is aligned with cell phone's longest side (pointing down). I consider a cell phone with three accelerometers (aligned along the phone's axes). If the cell phone is static or moving with constant velocity, each accelerometer measures the component of the gravity vector along its axis. This allows one to estimate the direction of gravity in the phone coordinate frame. The roll and pitch of the phone with respect to the natural frame can be estimated using the following equations:

$$\begin{aligned} \beta &= \arcsin\left(-a_x/g\right) \\ \alpha &= \arcsin\left(a_y/g \cdot \cos\beta\right) \end{aligned} \tag{1.1}$$

where $a_x$ and $a_y$ are the measurements from the $X$-axis and $Y$-axis accelerometer separately. $g$ is the gravity acceleration($g = 9.8m/s^2$). $\alpha$ is the roll angle, the rotation angle around X axis, while $\beta$ is the pitch angle, the rotation angle around Y axis.

In order to compute the yaw (the rotation angle around Z axis), one needs the help of either the digital compass or the onboard gyroscopes. Unfortunately, the magnetic field cannot be assumed to be constant (due to external sources) so the compass does not provide reliable measurements. The three gyroscopes, aligned with the three

axes of the phone, measure the angular velocity around 3 axes of the phone coordinate frame. The angular velocity with respect to the inertial natural frame can be represented by the following equations [7]:

$$\begin{aligned} \dot{\alpha(t)} &= \omega_y \cos\beta(t) - \omega_x \sin\beta(t) \\ \dot{\beta(t)} &= \omega_z + (\omega_y \sin\beta(t) + \omega_x \cos\beta(t)) \tan\alpha(t) \\ \dot{\gamma(t)} &= (\omega_y \sin\beta(t) + \omega_x \cos\beta(t)) \,/\, \cos\alpha(t) \end{aligned} \tag{1.2}$$

where $\alpha(t)$, $\beta(t)$ and $\gamma(t)$ are roll, pitch and yaw angles of the rotation of cell phone in natural frame at a given time $t$. Integrating the angular velocity in equation (1.2) gives the rotation angle. Integration, however, accumulates noise and may lead to drift. Therefore some research fuses the accelerometer measurements and gyroscope measurements together for robust orientation estimation, e.g. using a Kalman filter [16, 11, 4].

In order to assess the accuracy of pose estimation as provided by the inertial sensors, it is important to first measure the quality of such sensors. Accordingly, my thesis research addressed a specific question: how can one assess the accuracy of the inertial sensors contained in a mobile device such as a cell phone? Knowledge of the precise characteristics of these sensors is critical for applications that require precise inertial measurements. Examples include the use of the accelerometer as an inclinometer (e.g. to support precise camera pointing) or as part of a dead reckoning system to track the pose of the cell phone. Unfortunately, accuracy data is often not available. Even when the model of the accelerometer mounted on a cell phone is disclosed, the data sheet for the sensor is not always comprehensive.

This article focuses on benchmarking the accelerometers in two commercially available smartphones (the Nokia N97 and the iPhone 4) and the gyroscopes in the iPhone 4. It is organized as follows: Ch.2 describes the cell phones considered in my experiments and introduces the experimental setting along with the procedure for misalignment estimation first. Ch.3 presents experimental results in terms of random noise, error bias, error bias correlations, and orientation reconstruction error. Ch.4 provides the conclusions, highlights the shortcomings of our procedure, and suggests alternative methods for accuracy benchmarking.

# Chapter 2

# Methodology

In benchmarking the accuracy of accelerometers, I introduce a procedure for the comparison of accelerometer data with ground truth values. The cell phone is attached to the tip of a robotic arm, whose position and orientation can be controlled with great accuracy. The misalignment between the onboard accelerometer's reference system and the arm's reference system is estimated using Horn's algorithm [12] from multiple measurements. I use this procedure to benchmark the accuracy of two cell phone models: an Apple iPhone 4, and a Nokia N97. In order to benchmark the gyroscopes in the iPhone 4, I used a similar procedure. The cell phone is attached to the tip of a robotic arm, whose rotation speed can be controlled with great accuracy. The misalignment between the onboard gyroscope's reference system and the arm's reference system is estimated using Horn's algorithm [12] from multiple measurements too.

## 2.1 Benchmarking the Accuracy of Accelerometers in Cell Phones

### 2.1.1 Devices

I have considered two different cell phone models in this work: an Apple iPhone 4 and a Nokia N97. Both cell phones have 3-axes MEMS accelerometers that are easily accessible through an API. The iPhone 4 uses the ST LIS331DLH accelerometer produced by ST Microelectronics. According to the provided data sheet [1], the ST LIS331DLH produces 12 bit data with controllable range (from ±2.0 g to ±8.0 g) and associated sensitivity (from ±1.0 milli-g/digit to ±3.9 milli-g/digit) and variable rate (from 50 Hz to 1000 Hz). In the iPhone 4 implementation, the range is set to ±2.0 g, while the reading rate is of 56 Hz (measured based on the time stamps associated with each measurement). According to the Symbian S60 C++ Sensor API documentation, the accelerometer onboard the N97 has range of ±2.0 g and produces 8 bit data. The reading rate is of 10 Hz.

In order to measure the accuracy of the accelerometer measurements, I attached each cell phone in turn to the tip of a DENSO VM-60B1 6-axis robotic arm via a customized support (see Fig. 2.1). With a nominal position repeatability of ±0.07 mm [2], the DENSO VM-60B1 can be safely assumed to provide accurate "ground truth" measurements of the position and attitude of the arm tip (defined with respect to a

[1] http://www.st.com/internet/com/TECHNICAL_RESOURCES/TECHNICAL_LITERATURE/DATASHEET/CD00213470.pdf

[2] http://www.aarobotics.co.uk/images/stories/assests/DENSO/Serie%20VM%201000-1300%20mm_GB.pdf


Figure 2.1: The experimental setting. Note the iPhone 4 attached to the DENSO robotic arm.

given reference system whose $Z$ axis is aligned with the direction of gravity).

The robotic arm's control system allows one to precisely specify the arm tip's orientation with respect to the arm's reference system. My experiments were conducted by rotating the arm tip holding the cell phone to a number of stationary positions, and keeping it still in each position for a few seconds while reading from the cell phone's accelerometers was taking place. The cell phone's reference system is defined by the three axes along which the accelerometers are oriented, assumed to be orthogonal to each other. If the cell phone's reference system is aligned with the arm tip's reference system, then the orientation of the cell phone with respect to the robotic arm's reference system is known at each time, along with the "ground truth" component of the gravity vector onto the three axes $\overrightarrow{a} = (a_x, a_y, a_z)$. The difference between the ground truth ($\overrightarrow{a}$) and the measured ($\overrightarrow{\hat{a}}$) acceleration values defines the error $\overrightarrow{e} = (e_x, e_y, e_z)$ for a

given orientation.

### 2.1.2 Solving for Misalignment

The cell phone's reference system is defined such that its $X$ axis is orthogonal to the screen (pointing towards the back of the phone), its $Y$ axis is aligned with the shortest side of the screen (pointing right), and its $Z$ axis is aligned with cell phone's longest side (pointing down). When placing each phone on the robotic arm, I tried to mechanically align these axes with the reference axes of the arm tip. However, a certain amount of error has to be expected, meaning that the orientation of the cell phone is only known up to a (constant but unknown) *misalignment* factor. In formulas:

$$\overrightarrow{a} = R^m R \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \tag{2.1}$$

where $R^m$ is the (unknown) misalignment rotation matrix, $R$ is the (known) rotation matrix relating the arm tip's orientation to the robot's reference system, and the entries of $\overrightarrow{a}$ are expressed in gravity acceleration units (g = 9.8 m/s$^2$).

In order to estimate $R^m$, I collected a large number of measurements of $\{\overrightarrow{a}_i\}$ for multiple (known) values of the rotation matrix $\{R_i\}$, and then used Horn's algorithm for closed-form absolute orientation estimation [12, 5]. Horn's method expresses the misalignment rotation matrix $R^m$ by a quaternion $\mathbf{r}^m$, and solves the following

minimization problem:

$$\mathbf{r}^m = \underset{||\mathbf{r}||=1}{\operatorname{argmin}} \sum_i \|\mathbf{q}_i - \mathbf{r}\mathbf{p}_i\mathbf{r}^*\|^2 \tag{2.2}$$

where $\mathbf{q}_i$ is the quaternion expression of the vector formed by the third column of $R_i$, $\mathbf{p}_i$ is the quaternion expression of $\overrightarrow{\hat{a}}_i$ (expressed in units of $g$), $\mathbf{r}^*$ is the conjugate complex of $\mathbf{r}$, and the products and norms in the formula are expressed in quaternion algebra. The solution of (2.2) is obtained by eigenvalue analysis. For more details, the reader is referred to Horn's original article [12].

Once the misalignment rotation represented by $R^m$ is computed, I rotate the measurement vector by $R_m^{-1}$ in order to obtain "compensated" (or "re-aligned") values.

### 2.1.3 Rotation Representation

I represent the rotation induced by the matrix $R$ by Euler axis and rotation [3], as shown in Fig. 2.2. Note that the rotation axis is characterized by two parameters (elevation and azimuth), and thus, in order to test all possible orientations of the cell phone, one would have to consider the whole range of elevation ($\phi$), azimuth ($\psi$), and rotation ($\theta$) angles. In fact, for my measurements, I can leave the azimuth angle fixed. This is because I am only interested in the 2-D space spanned by the (constant modulus) gravity vector as seen in the reference system of the cell phone. Any possible value of $\overrightarrow{a}$ can be obtained by a rotating the cell phone by a certain rotation angle around a certain axis with fixed azimuth.

[3] I note in passing that the quaternion representation of the misalignment rotation $\mathbf{r}^m$ is readily translated into Euler axis and rotation, making this representation particularly efficient.

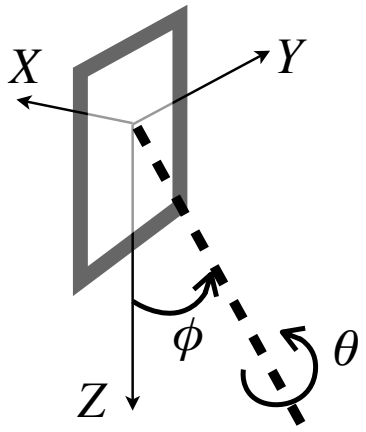


Figure 2.2: The chosen representation of rotation by rotation axis (with elevation $\phi$) and rotation angle $\theta$. The grey parallelogram represents the shape of the cell phone.

In the experiments, I sampled the elevation angle $\phi$ uniformly between -80° and 80° (17 samples). For each elevation angle, I sampled the rotation angle $\theta$ uniformly between -170° and 170° (35 samples). At each location, the arm was kept still for three seconds; however, data was not acquired during the first and third second, to avoid the risk of residual motion.

## 2.2 Benchmarking the Accuracy of Gyroscopes in Cell Phones

### 2.2.1 Devices

The Apple iPhone 4 has 3-axes MEMS gyroscopes that are easily accessed through an API. The iPhone 4 uses the ST L3G4200D angular rate sensor produced

by ST Microelectronics. According to the provided data sheet [4], the ST L3G4200D produces 16 bit data with controllable range from ±250 *degrees/sec* to ±2000 *degrees/sec* and associated sensitivity and controllable output rate (100/200/400/800Hz). In the iPhone 4 implementation, the range is set to ±2000 *degree/sec*, while the reading rate is of 56 Hz (measured based on the time stamps associated with each measurement).

My experiments were conducted by rotating the robotic arm tip holding the cell phone at different speed about different stationary axes while reading from the cell phone's gyroscopes. The cell phone's reference system is defined by the three axes along which the gyroscopes are oriented, assumed to be orthogonal to each other. If the cell phone's reference system is aligned with the arm tip's reference system, then rotation velocity and orientation of rotation axis of the cell phone with respect to the robotic arm's reference system is known at each time, along with the "ground truth" components of angular vector onto the three axes $\overrightarrow{\omega} = (\omega_x, \omega_y, \omega_z)$. The difference between the ground truth ($\overrightarrow{\omega}$) and the measured ($\overrightarrow{\hat{\omega}}$) angular velocity vector defines the error $\overrightarrow{e} = (e_x, e_y, e_z)$ for a given rotation axis $(\phi, \psi)$ and the magnitude of the rotation velocity $v \equiv ||\overrightarrow{\omega}||$.

### 2.2.2 Solving for Misalignment

Similar to the method described in section 2.1.2, if I define $R$ as the direction cosine matrix from the robot frame to phone frame, the objective is to find the optimal

[4] `http://www.st.com/internet/com/TECHNICAL_RESOURCES/TECHNICAL_LITERATURE/DATASHEET/CD00265057.pdf`

solution of $R$ such that the error is minimum.

$$\hat{R} = \underset{R}{argmin} \sum_{i=1}^{N} \left\| \vec{\omega}_i' - R\,\vec{\omega}_i \right\|^2 \tag{2.3}$$

where $\vec{\omega}_i$ and $\vec{\omega}_i'$ represent the controlled and measured angular velocity vector seperately in the $i^{th}$ experiment, $i = 1, 2, ..., N$. Horn's method [12] expresses Eq.2.3 with quaternions:

$$\hat{\mathbf{r}} = \underset{\mathbf{r}}{argmin} \sum_{i=1}^{N} \left\| \mathbf{q}_i' - \mathbf{r}\,\mathbf{q}_i\,\mathbf{r}^* \right\|^2 \tag{2.4}$$

where $\mathbf{q_i}$ and $\mathbf{q_i}'$ are the quaternion expression of the 3D vector $\vec{\omega}_i$ and $\vec{\omega}_i'$ respectively, $\mathbf{r}$ is the quaternion expression of the misalignment rotation matrix $R$. The solution of (2.4) is obtained by eigenvalue analysis. For more details, the reader is referred to Horn's original article [12].

Once the misalignment rotation represented by $\hat{R}$ is computed, I rotate the measurement vector by $\hat{R}^{-1}$ in order to obtain "compensated" (or "re-aligned") values.

# Chapter 3

# Experimental Results

## 3.1 Benchmarking the Accuracy of Accelerometers

### 3.1.1 Noise (Random Error)

By *noise* ($n$) I mean the difference between the measured acceleration values and their average during the acquisition period. More precisely: the noise $n(\phi, \theta, k)$ for the $k$-th sample during the acquisition period for elevation $\phi$ and rotation $\theta$ is defined by:

$$n(\phi, \theta, k) = \hat{a}(\phi, \theta, k) - \bar{a}(\phi, \theta) \tag{3.1}$$

where $\bar{a}(\phi, \theta)$ is the mean of the measured values within the acquisition period. The variance of the measured data is thus defined as the mean of $n^2(\phi, \theta, k)$ over all elevation and rotation angles and over all samples in each acquisition period. The standard deviation is, as always, the squared root of the variance. Noise is due to various cases (e.g., thermal noise), compounded by quantization effects.

| | $\sigma_x$ | $\sigma_y$ | $\sigma_z$ |
|---|---|---|---|
| iPhone 4 | 0.4 | 0.5 | 0.3 |
| N97 | 9.6 | 10.5 | 6.8 |

Table 3.1: The standard deviation of noise over the three accelerometer axes for the iPhone 4 and the N97 cell phones. All units are in *milli-g* $=10^{-3} \cdot 9.8\,\mathrm{m/s^2}$.

Tab. 3.1 shows the measured standard deviation of the noise over the three accelerometer axes for both cell phones considered. These measurements are consistent with the expected quantization noise, which, as well known [10], has standard deviation equal to $\Delta/\sqrt{12}$, where $\Delta$ is the quantization step. Assuming that the data is uniformly quantized between -2 g and 2 g for both cell phones, the quantization interval $\Delta$ is equal to 4 g$/2^{12} = 1$ milli-g for the iPhone 4 and 4 g$/2^8 = 16$ milli-g for the N97, resulting in expected noise standard deviations of 0.3 milli-g for the iPhone 4 and of 4.5 milli-g for the N97. The measured noise standard deviation is shown to be no larger than approximately twice the quantization noise standard deviation.

### 3.1.2 Bias

The *bias* ($\bar{e}$) is defined the difference between the mean measured value within an acquisition period and its value as predicted by knowledge of the actual orientation of the camera. In formulas:

$$\bar{e}(\phi, \theta) = \bar{a}(\phi, \theta) - a(\phi, \theta) \tag{3.2}$$

The bias thus represents the systematic measurement error of the sensors.

Fig. 3.1 shows, for each accelerometer axis, the histogram of the bias $\bar{e}$. It can be noticed that mean of the bias (the *average bias*) is noticeably non-null in some cases

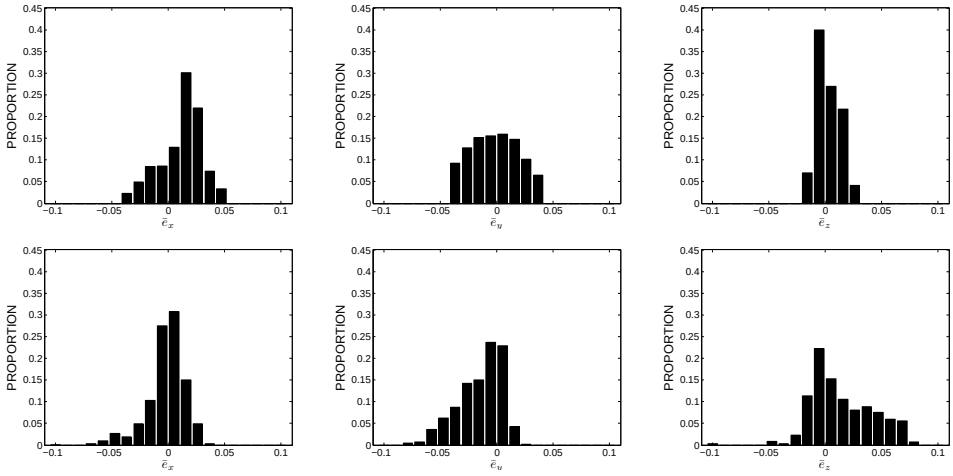


Figure 3.1: Histograms of the bias $\bar{e}$ for the iPhone 4 (Top) and for the N97 (Bottom). Left column: $X$ axis. Middle column: $Y$ axis. Right column: $Z$ axis. All units are in $g = 9.8\,\mathrm{m/s^2}$.

| | $e_{x,\mathrm{RMS}}$ | $e_{y,\mathrm{RMS}}$ | $e_{z,\mathrm{RMS}}$ |
|---|---|---|---|
| | *before bias average subtraction* | | |
| iPhone 4 | 22.5 | 21.1 | 9.7 |
| N97 | 17.4 | 25.9 | 21.9 |
| | *after bias average subtraction* | | |
| iPhone 4 | 18.8 | 20.1 | 9.5 |
| N97 | 17.3 | 20.2 | 21.2 |

Table 3.2: The root mean square error of the bias $\bar{e}$ for the three accelerometer axes for the iPhone 4 and the N97 cell phones. All units are in *milli-g* $=10^{-3} \cdot 9.8\,\mathrm{m/s^2}$.

(in particular, for the $X$ axis for the iPhone, and for the $Y$ axis for the N97). In fact, the average bias could be removed by simply subtracting it from the measured signal. Tab. 3.2 shows the root mean square (RMS) of the bias for each axis before and after average bias subtraction. Note in passing that the bias RMS computed on the measured data *before* re-alignment (as described in Sec. 2.1.2) is about 30 times (or more) higher, providing evidence that the re-alignment procedure is effective.

Fig. 3.2 represents the bias $\bar{e}$ as a function of $(\phi, \theta)$. It is interesting to note

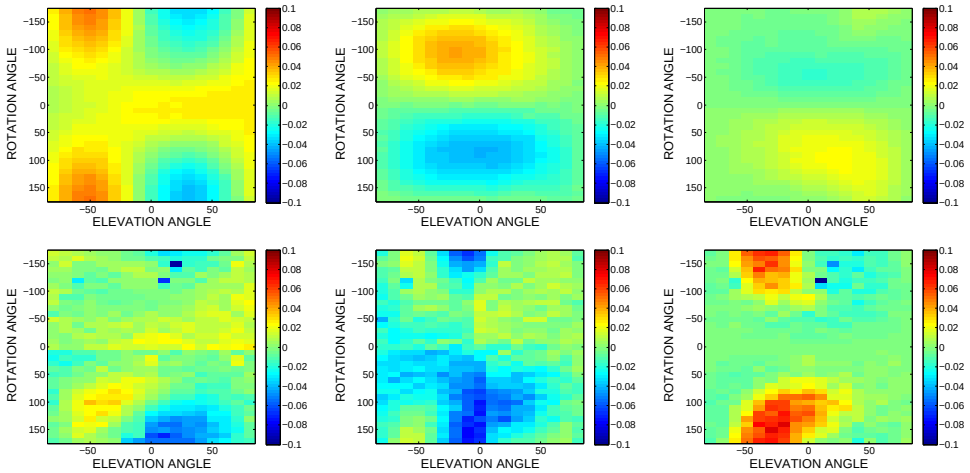


Figure 3.2: Representation of the bias $\bar{e}$ as a function of elevation $\phi$ and rotation $\theta$ (in *degree*) for the iPhone 4 (top) and the N97 (bottom). Left column: $X$ axis. Middle column: $Y$ axis. Right column: $Z$ axis. The bias is represented in units of $g = 9.8\,\mathrm{m/s^2}$.

that the bias is well correlated to the cell phone orientation. The fact that the bias is more "discontinuous" (as a function of $\phi$ and $\theta$) for the N97 than for the iPhone 4 could be imputed to the N97's larger noise variance, combined with the smaller number of samples per acquisition interval considered, which concur to higher variance in the estimation of the average values $\bar{a}(\phi, \theta)$ and thus of the bias.

A scattergram of bias $\bar{e}$ vs. acceleration $a$ for each axis is shown in Fig. 3.3. Two observations can be drawn from these plots. Firstly, there seems to be only a weak correlation between bias $\bar{e}$ and the gravity value $a$ at each axis. For example, the bias in the $Z$ axis of the iPhone 4 has higher variance for gravity values that are close to 0, while the bias in the $Z$ axis of the iPhone 4 has higher variance for gravity values close to $g$. Secondly, the point cloud for the iPhone 4 is characterized by a highly correlated pattern (a set of curves), rather than being randomly distributed as for the N97. I

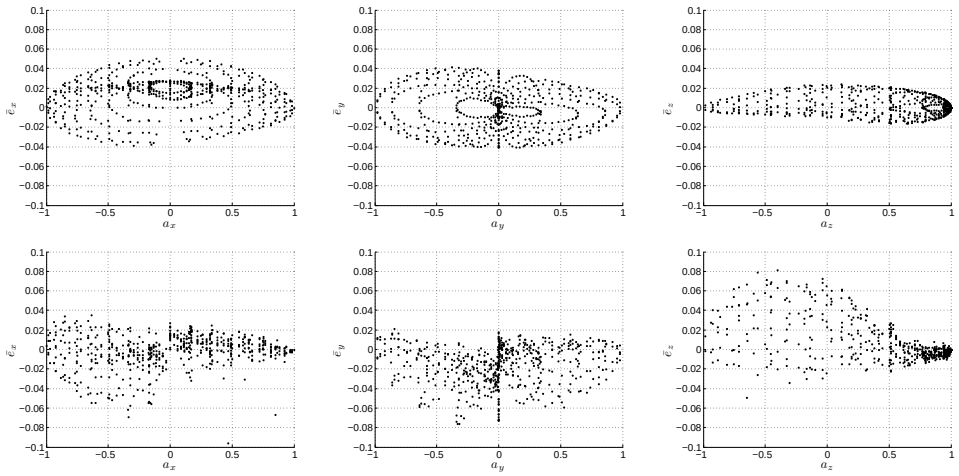

Figure 3.3: Bias vs. acceleration for the iPhone 4 (top) and the N97 (bottom). Left column: $X$ axis. Middle column: $Y$ axis. Right column: $Z$ axis. All units are in $g = 9.8\,\mathrm{m/s^2}$.

conjecture that the vector $(a, \bar{e})$ moves on a smooth curve as the cell phone is being rotated, and that different curves correspond to different elevations.

Fig. 3.4 shows scattergrams of bias relative to one axis against bias relative to another axis. These plots are meant to highlight any correlation between axes. Indeed, it is seen that, for the iPhone 4, the bias on the $Y$ axis correlates negatively with the bias on the $Z$ axis.

### 3.1.3 Orientation Estimation Error

The measured acceleration values may be used to estimate the orientation of the cell phone, under the assumption that the azimuth is fixed. It should be noted that the relation between the acceleration $\overrightarrow{a}$ and the orientation (expressed in Euler axis and rotation) is highly non-linear, hence the same bias error may result in larger or

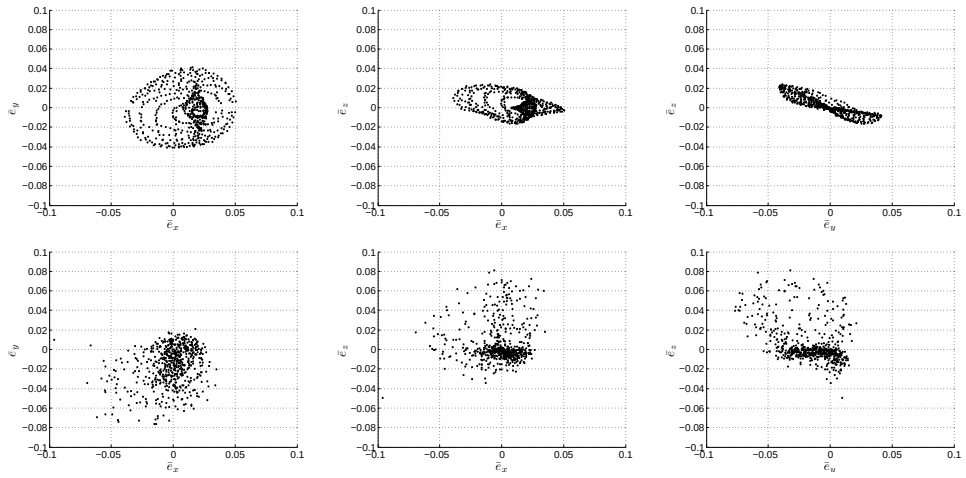

Figure 3.4: Bias on one axis vs. bias on another axis for the iPhone 4 (top) and the N97 (bottom). Left column: $X$ axis. Middle column: $Y$ axis. Right column: $Z$ axis. All units are in $g = 9.8\,\mathrm{m/s^2}$.

smaller error on the estimated orientation for a given orientation. Note also that when the rotation angle $\theta$ is equal to $0°$, the elevation angle $\phi$ of the rotating axis cannot be estimated. Likewise, the rotation angle $\theta$ cannot be estimated when the elevation angle $\phi$ is equal to $\pm 90°$. As shown in Fig. 3.5, within the range of orientations considered in our example, the error in elevation estimation ranges between $-1.8°$ and $3.4°$ for the iPhone 4 and between $-1.5°$ and $5.2°$ for the N97, while the error in rotation estimation ranges between $-11.5°$ and $8.7°$ for the iPhone 4 and $-14.9°$ and between $15.2°$ for the N97.

In Fig. 3.6, I show the residual rotation and elevation angle after rotating back by the inverse of the rotation estimated by the sensors. In the ideal case of correct acceleration measurement, the rotation would be estimated perfectly, and the residual rotation would be null. Note in passing that Fig. 3.6 provides indirect evidence that

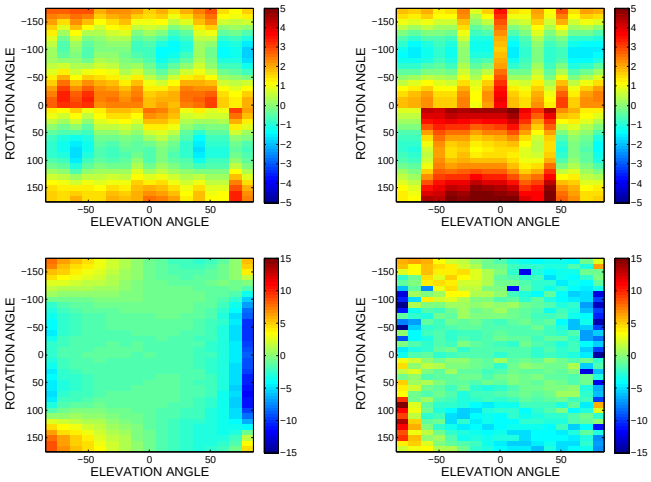


Figure 3.5: Representation of the orientation estimation error as a function of elevation $\phi$ and rotation $\theta$ for the iPhone 4 (left) and the N97 (right). Top: error in the estimation of the elevation angle $\phi$. Bottom: error in the estimation of the rotation angle $\theta$. All units are in *degree*.

the dependence of the error bias on the device's orientation should not be attributed to errors in the estimation of the misalignment rotation $R_m$. If this were the case, the residual would be constant for all orientations - which, as seen in Fig. 3.6, is not true in my case.

## 3.2 Benchmarking the Accuracy of Gyroscopes

In the experiments of benchmarking the accuracy of gyroscopes, I sampled both the elevation angle $\phi$ and the azimuth angle $\psi$ of the rotation axis in $\{0, \pm 30, \pm 60, \pm 90\}$ *degrees*. For each axis, the arm was rotated at four different speed $\{13, 46, 78, 111\}$ *degrees/sec*. Data was not acquired during the first and last 3 seconds, in order to allow the system to reach a stationary velocity.

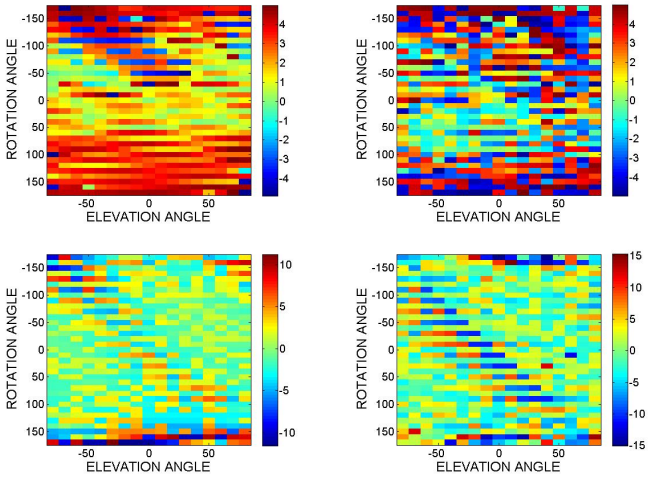


Figure 3.6: Representation of the residual elevation and rotation after rotating back according to the inverse of the estimated rotation (based on sensor data), as a function of elevation $\phi$ and rotation $\theta$ for the iPhone 4 (left) and the N97 (right). Top: residual elevation angle $\phi$. Bottom:residual rotation angle $\theta$. All units are in *degree*.

| | $\sigma_x$ | $\sigma_y$ | $\sigma_z$ |
|---|---|---|---|
| iPhone 4 | 0.3 | 0.3 | 0.4 |

Table 3.3: The standard deviation of noise over the three gyroscope axes for the iPhone 4. All units are in *degree/sec*.

### 3.2.1 Noise (Random Error)

Tab. 3.3 shows the measured standard deviation of the noise over the three gyroscope axes for iPhone 4. These measurements are consistent with the expected quantization noise, which, using the methods in section.3, is $4000/2^{16}/\sqrt{12} = 0.0176$ *degrees/sec*. The measured noise standard deviation is shown to be no larger than approximately twice the quantization noise standard deviation.

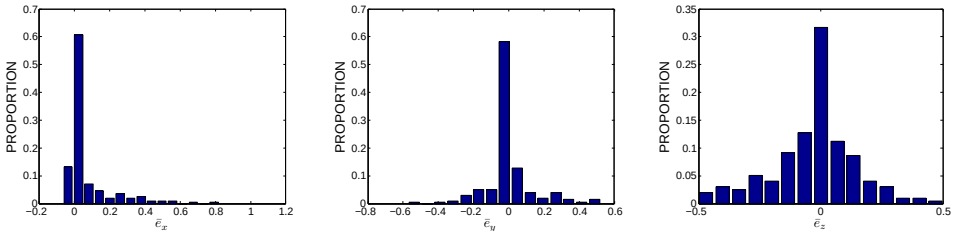


Figure 3.7: Histograms of the bias $\bar{e}$ for the iPhone 4. Left: $X$ axis. Middle: $Y$ axis. Right: $Z$ axis. All units are in *degree/sec*.

| | $e_{x,\mathrm{RMS}}$ | $e_{y,\mathrm{RMS}}$ | $e_{z,\mathrm{RMS}}$ |
|---|---|---|---|
| | *before bias average subtraction* | | |
| iPhone 4 | 0.15 | 0.13 | 0.17 |
| | *after bias average subtraction* | | |
| iPhone 4 | 0.14 | 0.12 | 0.17 |

Table 3.4: The root mean square error of the bias $\bar{e}$ for the three gyroscope axes for the iPhone 4. All units are in *degree/sec*.

### 3.2.2 Bias

Fig. 3.7 shows, for each gyroscope axis, the histogram of the bias $\bar{e}$. Tab. 3.4 shows the root mean square (RMS) of the bias for each axis before and after average bias subtraction. Note that the bias RMS computed on the measured data *before* re-alignment (as described in Sec. 2.2.2) is about 10 times higher, providing evidence that the re-alignment procedure is effective.

# Chapter 4

# Conclusions

This contribution has introduced a methodology for precise benchmarking of the inertial sensors within a mobile device. Use of a precisely controllable robotic arm allows for repeatable, accurate ground truth measurements. Even if a robotic arm is not available, one could still perform similar measurements by attaching a high-quality sensor to the device, and comparing the readings of the two sensors. Even in this case, the misalignment between the external, high-quality sensor and the internal sensor being assessed need to be estimated, for example using the procedure proposed in Sec. 2.1.2 and Sec. 2.2.2.

The proposed benchmarking method in this thesis, while accurate, has a number of shortcomings. Firstly, it can only measure the accuracy of the device's accelerometers within $\pm 1$ g. Although for applications that require estimating the inclination of the cell phone this range of measurements is sufficient, other applications may require measurement of larger or smaller accelerations. This could be achieved by moving the

device in high-acceleration trajectories (e.g., shaking it, or perhaps mounting it on a rotating arm). However, in this case, it is critical that the measurements from the internal and the external high-quality accelerometer (whether attached to the device, or in the robotic arm holding the device) should be time-aligned, for example by means of time stamps and clock synchronization. We were able to avoid this procedure by only taking measurements when the cell phone was in stable positions, as explained in Sec. 2.1.3.

The experimental data with accelerometers in Apple iPhone 4 and the Nokia N97 has highlighted the presence of a non-negligible bias that depends on the orientation. This bias can be as high as 0.1 g. Subtracting the mean component of the bias from the measurements provides a slight improvement in terms of RMS. The bias error induces an error of up to several degrees when estimating the cell phone's orientation based on data from the accelerometer that is slightly higher for the N97 than for the iPhone 4. In addition to bias, there is a random noise component that is markedly higher for the N97 (up to 10 milli-g) than for the iPhone 4, consistent with the fact that the data from the iPhone 4 is quantized at 12 bits/sample vs. 8 bits/sample for the N97.

More work is needed to understand why the bias $\bar{e}$ correlates with the cell phone's orientation (see Fig. 3.2) while only loosely correlating with the actual acceleration (see Fig. 3.3). One may also conjecture that the accelerometer axes may not be precisely orthogonal to each other. This could also explain the correlation between the bias on the $Y$ and $Z$ axes for the iPhone 4 that is clearly noticeable in Fig 3.4.

# Bibliography

[1] Knfb reading technology. *http://www.knfbreader.com*.

[2] Sandip Agrawal, Ionut Constandache, Shravan Gaonkar, Romit Roy Choudhury, Kevin Caves, and Frank DeRuyter. Using mobile phones to write in air. *Proceedings of the 9th international conference on Mobile systems, applications, and services*, pages 15–28, 2011.

[3] Michael Aron, Gilles Simon, and Marie-Odile Berger. Use of inertial sensors to support video tracking: Research articles. *Comput. Animat. Virtual Worlds*, 18:57–68, February 2007.

[4] Gabriele Bleser and Didier Stricker. Advanced tracking through efficient image processing and visual-inertial sensor fusion. *Computers and Graphics*, 33(1):59–72, 2009.

[5] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[6] J. Coughlan and R. Manduchi. Functional assessment of a camera phone-based

wayfinding system operated by blind and visually impaired users. *International Journal on Artificial Intelligence Tools, Special Issue on Artificial Intelligence Based Assistive Technologies: Methods and Systems for People with Disabilities*, 18(3), 2009.

[7] E. Eyjolfsdottir and M. Turk. Multisensory embedded pose estimation. *Applications of Computer Vision (WACV), 2011 IEEE Workshop on*, pages 23 –30, Jan. 2011.

[8] D. Gaida, A. Stuhlsatz, and H. G. Meier. Fusion of visual and inertial measurements for pose estimation. *Proceedings of the workshop on Emotion, Speech And Face Recognition with Advanced Classifiers*, 2008.

[9] S. Gammeter, A. Gassmann, L. Bossard, T. Quack, and L.J. van Gool. Server-side object recognition and client-side object tracking for mobile augmented reality. *Proc. International Workshop on Mobilef Vision*, 2010.

[10] A. Gersho and R. Gray. *Vector quantization and signal compression.* Kluwer Academic Publishers, Norwell, MA, USA, 1991.

[11] Jeroen D. Hol, Thomas B. Schön, Henk Luinge, Per J. Slycke, and Fredrik Gustafsson. Robust real-time tracking by fusing measurements from inertial and vision sensors. *J. Real-Time Image Processing*, 2(2-3):149–160, 2007.

[12] B.K.P. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America*, 4:629–642, 1987.

[13] M Hwangbo, J S Kim, and T Kanade. Gyro-aided feature tracking for a moving camera: fusion, auto-calibration and gpu implementation. *The International Journal of Robotics Research*, 30(14):1755–1774, Dec. 2011.

[14] Georg Klein and Tom Drummond. Tightly integrated sensor fusion for robust visual tracking. *Proc. British Machine Vision Conference (BMVC'02)*, pages 787–796, 2002.

[15] Jiayang Liu, Lin Zhong, Jehan Wickramasuriya, and Venu Vasudevan. *uWave: Accelerometer-based personalized gesture recognition and its applications*, volume 5. Elsevier Science Publishers B. V., Amsterdam, The Netherlands, Dec. 2009.

[16] J. Vaganay, M. J. Aldon, and A. Fournier. Mobile robot attitude estimation by fusion of inertial data. *IEEE International Conference on Robotics and Automation*, 1:277–282, May 1993.

[17] D. Wagne, G. Reitmayr, A. Mulloni, T. Drummond, and D. Schmalstieg. Pose tracking from natural features on mobile phones. *Proceedings of the 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*, pages 125–134, 2008.

[18] Oliver Woodman. An introduction to inertial navigation. *American Journal of Physics*, 77(UCAM-CL-TR-696):844, Sep. 2007.

[19] Qingxuan Yang, Chengu Wang, Yuan Gao, Hang Qu, and Edward Y. Chang. Inertial sensors aided image alignment and stitching for panorama on mobile phones.

In *Proceedings of the 1st international workshop on Mobile location-based service*, MLBS '11, pages 21–30, New York, NY, USA, 2011. ACM.