

CS60055: Ubiquitous Computing

Location, Gesture and Activity Sensing

Part I: The Fundamentals of Motion Tracking

Reading Material

Sandip Chakraborty

sandipc@cse.iitkgp.ac.in

Department of Computer Science and Engineering

Indian Institute of Technology Kharagpur

Contents

1	Introduction: Sensing in the Age of Wearables	2
2	Inertial Measurement Units: An Overview	2
3	Reference Frames and Proper Acceleration	3
4	The Equivalence Principle	5
5	Effect of Device Rotation on the Gravity Offset	6
6	Representing Orientation: Euler Angles	7
7	The Direction Cosine Matrix	8
8	Estimating Roll and Pitch from Accelerometer Data	9
9	Measuring Yaw: The Role of the Gyroscope	11
10	Sensor Fusion for Robust Orientation Estimation	11
11	Looking Ahead: Gyroscopic Effects and Self-Balancing	13

1 Introduction: Sensing in the Age of Wearables

Wearable computing has moved from a niche research pursuit to an everyday reality. Devices such as GPS trackers, smart glasses, fitness bands, smartwatches and smart rings are now commonplace, embedding an increasing amount of sensing capability directly onto the human body.



Figure 1: A collection of common wearable devices: GPS tracker, smart glasses, fitness band with heart rate display, smart ring, smartwatch and glucometer.

Wearable form factors span virtually the entire body, from smart glasses and smartwatches at the head and wrist, through smart shirts and belts on the torso, down to smart socks and shoes at the feet, with several other form factors (bracelets, finger worn devices, rings, pants) occupying the space in between.

These devices are deployed across a broad range of application areas. In healthcare, wearables support continuous monitoring of vital signs and chronic conditions, for example cardiac monitoring, glucose monitoring or blood pressure tracking. In what is often termed secondary healthcare or general wellness, they support sleep tracking and everyday health monitoring, and in fitness and sports, they support activity tracking, performance monitoring and recovery.

To support this range of applications, wearable devices integrate a remarkably diverse palette of sensing modalities. These include, among many others, inertial sensors (IMU), altimeters, bio-acoustic sensors, blood pressure sensors, ambient brightness sensors, cameras, electrocardiogram (ECG) sensors, electrodermal activity (EDA) sensors, electromyogram (EMG) sensors, fiber optic sensors, compasses, GPS receivers, galvanic skin response (GSR) sensors, humidity sensors, magnetometers, infrared proximity and temperature sensors, LEDs, glucometers, pressure sensors, microphones, piezoelectric sensors, RFID tags, spectrometers, thermometers, textile sensors, ultrasound sensors, weight sensors, optical sensors, mmWave radar sensors, electrodes, inductive wire sensors and infrared photo-reflective sensors.

Of all these modalities, the Inertial Measurement Unit (IMU), comprising an accelerometer and a gyroscope (and often a magnetometer), plays a particularly foundational role. It is present in virtually every wearable device and underlies the vast majority of location, gesture and activity sensing applications built on wearable platforms. This part of the course develops the fundamentals of motion tracking using IMU sensing, starting from first principles.

To motivate the discussion concretely, consider a deceptively simple question: **given the accelerometer and gyroscope data recorded by a smartwatch worn during a period of physical activity, how would one distinguish between the wearer walking versus running?**

At first glance this may appear to be a simple matter of examining the magnitude or the frequency of the recorded acceleration signal. As we shall see over the course of this material, correctly answering even this seemingly elementary question requires a precise understanding of what an IMU actually measures, how gravity, orientation and genuine dynamic motion interact within the raw sensor readings, and how these effects can be systematically disentangled from one another.

2 Inertial Measurement Units: An Overview

An Inertial Measurement Unit (IMU) is an electronic sensing package composed of an accelerometer, a gyroscope, and, in many implementations, a magnetometer. Collectively these sensors allow an IMU to measure three related



Figure 2: Wearables from head to toe: smart glasses, smart watch, smart shirt, GPS/GPRS body control, Bluetooth key tracker, smart shoes, smart socks, smart pants, smart belt, smart ring, smart finger and smart bracelet.

but distinct physical quantities: the specific force, or proper acceleration, experienced by the device; its angular rate, that is, the rotation speed about each of its three axes; and, through appropriate processing of these two signals, its orientation in three dimensional space.

At the most basic level, an IMU detects linear acceleration and angular (rotation) rate. Beyond this, and of particular relevance to wearable computing, IMUs are very widely used to estimate the orientation of the device on which they are mounted. Orientation sensing is, in fact, one of the primary use cases for IMUs across the entire wearable device landscape. More broadly, IMUs are the workhorse sensor for maneuvering and navigation of any object through three dimensional space, whether that object is a smartphone, a wearable device, a drone, an aircraft or a spacecraft.

The remainder of this material develops, from first principles, precisely how an accelerometer and a gyroscope, used first individually and then jointly, allow us to estimate a device's orientation, and why understanding this process is essential for correctly interpreting IMU data in gesture and activity sensing applications.

3 Reference Frames and Proper Acceleration

To understand what an accelerometer actually measures, we must be precise about the notion of a reference frame, and about the particular type of acceleration that inertial sensors report.

A **rest frame** is the coordinate system in which a given particle or object is, at a particular instant, at rest, that is, has zero velocity. Every moving object has its own instantaneous rest frame, even though that frame may itself be accelerating with respect to some external observer.

The quantity that an accelerometer physically measures is called the **proper acceleration**: the rate of change of velocity of a body, measured within its own instantaneous rest frame. Proper acceleration is the acceleration that would be measured by an observer who is in free fall, that is, an observer situated in a local inertial frame of reference. This is an important, and somewhat counter-intuitive, distinction: proper acceleration is in general *not* the same as coordinate acceleration, the rate of change of velocity measured with respect to some fixed external,




Figure 3: Range of wearable devices across the healthcare, secondary healthcare, and fitness/sports segments

for instance Earth based, coordinate system.

To build intuition, consider a simple question: **what acceleration value would be reported by an accelerometer resting on a table in a classroom?** Since the device and the table are stationary with respect to the room, one might naively expect a reading of zero along every axis.




Figure 4: A gyroscope shaped accelerometer resting on a classroom table, with the two horizontal axes reading 0 m/s² and the vertical axis reading 9.8 m/s²

The correct answer, however, is that the accelerometer reads approximately 9.8 m/s^2 along the vertical direction, that is, pointing away from the table surface, while the two horizontal axes correctly read 0 m/s^2 . The reason for this seemingly counter-intuitive result is that the table exerts a normal, or contact, force on the accelerometer that exactly balances the pull of gravity, preventing the device from free falling. Because the accelerometer is thereby held out of its natural, free falling, inertial trajectory, it registers a proper acceleration equal in magnitude to the local gravitational acceleration g , directed upward, opposite to the direction in which gravity pulls. This observation can be summarised concisely: *a free-falling accelerometer will read zero acceleration*, precisely because free fall is what it means to be in an inertial, or rest, frame with no contact forces acting on the body.

This same reasoning extends to any object resting on a surface. A ball lying at rest on the ground is subject to the identical argument: although we describe the ball informally as being at rest, an accelerometer attached to it

would register a reading of g directed upward, on account of the normal force exerted by the ground.

This raises a natural and important question: can we represent the motion of every object using this simple framework of rest frames and proper acceleration alone? As we shall see, extending this idea to devices that are simultaneously moving and rotating, as is the case for essentially all wearable devices, requires considerably more care, and forms the subject of the remainder of this material.

To sharpen the motivation further, consider a wearable accelerometer embedded in a pair of headphones worn by a person who is listening to music, for instance while seated on an aircraft and gently moving their head to the rhythm. Could a single accelerometer reading, by itself, allow us to conclude that the wearer is enjoying the music, purely on the basis of proper acceleration and orientation? This example foreshadows a central theme of wearable sensing: raw inertial data conflates the static effect of gravity, which encodes information about the device's orientation, with the dynamic linear acceleration produced by genuine movement. Disentangling these two components, and exploiting each of them appropriately, underlies virtually every gesture and activity recognition pipeline built on top of IMU data.

4 The Equivalence Principle

The physical basis for why a stationary accelerometer registers a non-zero reading equal to g lies in one of the foundational ideas of modern physics: Einstein's **equivalence principle**.

Einstein's original insight was that the empirically observed equality between *inertial mass*, the mass that resists acceleration in Newton's second law $F = ma$, and *gravitational mass*, the mass that determines the strength of gravitational attraction in Newton's law of gravitation, has a profound consequence: the acceleration imparted to any object by a gravitational field is completely independent of that object's material composition or any other of its properties. Writing out Newton's equation of motion for a body in a gravitational field, we have that (inertial mass) times (acceleration) equals (the intensity of the gravitational field) times (gravitational mass); it is only when the numerical values of the inertial and the gravitational mass coincide that the resulting acceleration becomes independent of the nature of the body. This universality of free fall is precisely the content of the equivalence principle.

One elegant consequence of the equivalence principle is that a falling object behaves in an identical manner whether it is released inside a rocket accelerating through empty space at a rate equal to the local gravitational acceleration g , or inside an equivalent rocket sitting at rest on a planet's surface.




Figure 5: A falling object behaves identically inside an accelerating rocket in free space (left) and at rest on a planet's surface (right)

An observer confined to either rocket, without access to any external reference, cannot distinguish between these two scenarios using purely local mechanical experiments.

In its simplified and most operationally useful form for our purposes, the equivalence principle states that *the effect of gravity on an object is indistinguishable from the effect of an equivalent acceleration*.

Consider a rigid body with a local coordinate frame attached to it, with axes X , Y and Z . In the absence of any external force this local frame constitutes what we call the **local inertial frame** of the object: the frame in which the object's own dynamics can be described using the ordinary laws of motion, without needing to invoke any fictitious forces.




Figure 6: A rigid body (a ball) with its own local coordinate axes X, Y and Z.




Figure 7: The same rigid body with the gravity vector g now acting along its local Z axis.

When gravity acts on the object, for instance because the object rests on the Earth's surface, the accelerometer attached to it reports its proper acceleration measured with respect to this local inertial frame. This measured value already includes the contribution due to gravity, an effect we refer to as the **gravity offset**. If the object also happens to be moving through space, undergoing genuine dynamic acceleration as would be the case for a wearable device during physical activity, the gravity offset must first be subtracted from the raw accelerometer reading before the actual dynamic acceleration experienced by the object, as it moves along the Earth's surface, can be recovered.

5 Effect of Device Rotation on the Gravity Offset

The preceding discussion implicitly assumed that the device's local axes are aligned in some fixed, known way relative to the direction of gravity. In practice, a wearable device such as a smartwatch is free to rotate arbitrarily on the wrist, and its local axes will therefore be oriented differently relative to gravity depending on how the device is worn, or on how the wrist happens to be held.

Suppose, for one particular orientation of the watch, the gravity offset happens to align with the device's local Y axis.

If the wrist is subsequently rotated, so that the watch now faces a different direction, the device's local axes rotate along with it, and consequently the gravity offset no longer aligns with the same axis as before.

This observation carries an important implication: the same physical quantity, the gravity offset, of constant magnitude g , will be distributed differently across the three accelerometer axes (x , y , z) depending on the device's orientation. To make use of the accelerometer for orientation sensing, or conversely, to correctly remove the gravity offset when we are interested only in the dynamic component of acceleration, we need a systematic method




Figure 8: A watch on the wrist with local axes X, Y, Z; the gravity offset points along the Y axis.




Figure 9: The watch shown in two different orientations. The local inertial frame changes with rotation, so the gravity offset (of fixed magnitude g) is distributed differently across the local X, Y and Z axes in the two cases: in the first, $Y = g$; in the second, $Y \neq g$.

for describing how the local frame is rotated relative to a fixed reference frame, and how a fixed vector such as the gravity offset transforms under that rotation.

This is precisely the role played by the Euler angle representation of orientation, together with the associated rotation, or Direction Cosine, matrices, which we develop next.

6 Representing Orientation: Euler Angles

Before specialising to roll, pitch and yaw, it is worth noting that the general problem of describing how an object's orientation updates under rotation is a well studied one, with several equivalent representations available, including rotation matrices, Euler angles and quaternions.

In this material we adopt the **Euler angle representation**, which is the most intuitive for the purposes of accelerometer and gyroscope based orientation sensing.

A rigid body moving freely in three dimensional space possesses six degrees of freedom: three describing its position, translation along the X, Y and Z axes, and three describing its orientation, rotation about each of these three axes.

The three rotational degrees of freedom are conventionally named **roll**, **pitch** and **yaw**, and together they are referred to as the Euler angles describing the orientation of an object with respect to a chosen frame of reference. Roll, denoted ϕ , is rotation about the forward (longitudinal) axis; pitch, denoted θ , is rotation about the lateral axis; and yaw, denoted ψ , is rotation about the vertical axis.

To build intuition for how these three rotations combine to describe a general orientation, consider a smart-watch, and imagine transforming its axes from an initial reference alignment to some final, tilted alignment through a sequence of three elementary rotations.

Step 1 (Roll). We first rotate the Y-Z plane by the roll angle ϕ , which has the effect of adjusting the direction of the Z axis through a rotation about the (unchanged) X axis.




Figure 10: The reference axes X, Y, Z of a watch alongside its rotated local axes x, y, z. We need to determine the components of the gravity offset along each of the rotated local axes.

As a consequence of this rotation, the Y axis is also displaced, it does not remain fixed; both the Y and Z axes acquire new positions, denoted Y' and Z', as a result of the roll rotation, while only the X axis remains unaffected.

Step 2 (Pitch). We next rotate the X-Z plane by the pitch angle θ , adjusting the X axis through a rotation about the Y axis.

Following this rotation, the Y axis, already displaced by the roll rotation to Y', requires further adjustment as well, since the pitch rotation is performed about the post-roll Y' axis and not about the original Y axis.

Step 3 (Yaw). Finally, we rotate the Y-X plane by the yaw angle ψ with respect to the (post-pitch) Z axis.

This sequential process, rotating first about one axis, then about a second, updated, axis, and finally about a third, twice-updated, axis, is the geometric procedure underlying any Euler angle representation of orientation. It is formalised using rotation matrices in the next section.

7 The Direction Cosine Matrix

Let R_1 denote a fixed reference frame, for instance the frame of an external observer or a level frame aligned with the local vertical, and let R_2 denote the local, or body, frame attached to the object whose orientation we wish to determine.

To describe the orientation of the object we must specify the rotation that transforms R_2 into, that is, aligns it with, R_1 .

By convention, all rotations are taken to follow the right hand rule and are treated as positive in the clockwise sense, when viewed along the positive direction of the rotation axis. An arbitrary orientation can always be constructed using a sequence of exactly three elementary rotations, each about one of the three coordinate axes, subject to the constraint that the second rotation must be about an axis different from the first, while the third rotation may be about either the same axis as the first, or the one remaining axis not yet used. Labelling the axes 1(= X), 2(= Y) and 3(= Z), this gives exactly twelve valid three-rotation sequences: 1-2-1, 1-3-1, 2-1-2, 2-3-2, 3-1-3, 3-2-3, 1-2-3, 1-3-2, 2-3-1, 2-1-3, 3-1-2 and 3-2-1. Among all twelve, the 3-2-1 sequence, corresponding to yaw followed by pitch followed by roll, is by far the most widely used convention, particularly in aircraft and vehicle navigation, and is the sequence we adopt for the remainder of this material.

Each elementary rotation, by an angle α about one of the three axes, is represented by one of the following rotation matrices:

$$M_1(\alpha) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha & \sin \alpha \\ 0 & -\sin \alpha & \cos \alpha \end{pmatrix}, \quad M_2(\alpha) = \begin{pmatrix} \cos \alpha & 0 & -\sin \alpha \\ 0 & 1 & 0 \\ \sin \alpha & 0 & \cos \alpha \end{pmatrix}, \quad M_3(\alpha) = \begin{pmatrix} \cos \alpha & \sin \alpha & 0 \\ -\sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (1)$$




Figure 11: The six degrees of freedom of a rigid body: translations (up/down, back/forward, left/right) and rotations (roll, pitch, yaw)

We substitute for α the appropriate value of roll (ϕ), pitch (θ) or yaw (ψ), depending on which physical rotation is represented at each stage of a given sequence. For a 3-2-1 sequence, for instance, the angle used in M_1 corresponds to roll, that used in M_2 to pitch, and that used in M_3 to yaw.

Having defined the three elementary rotation matrices, the complete rotation describing a given Euler angle sequence, known as the **Direction Cosine Matrix** (DCM), is obtained by multiplying the corresponding elementary matrices together, applied right to left in the order in which the rotations are physically carried out. For a 2-3-1 sequence, for example, the DCM is

$$B = M_1(\alpha_3)M_3(\alpha_2)M_2(\alpha_1). \quad (2)$$

Following the same convention, the DCM for the 3-2-1 sequence, yaw (ψ) followed by pitch (θ) followed by roll (ϕ), is $B = M_1(\phi)M_2(\theta)M_3(\psi)$, which, on carrying out the matrix multiplication, gives

$$B = \begin{pmatrix} \cos \theta \cos \psi & \cos \theta \sin \psi & -\sin \theta \\ \sin \phi \sin \theta \cos \psi - \cos \phi \sin \psi & \sin \phi \sin \theta \sin \psi + \cos \phi \cos \psi & \sin \phi \cos \theta \\ \cos \phi \sin \theta \cos \psi + \sin \phi \sin \psi & \cos \phi \sin \theta \sin \psi - \sin \phi \cos \psi & \cos \phi \cos \theta \end{pmatrix} \quad (3)$$

This matrix B transforms the components of any fixed vector, expressed in the reference frame R_1 , into its corresponding components in the rotated body frame R_2 . Equivalently, B encodes the complete orientation of the body frame relative to the reference frame using just the three Euler angles ϕ , θ and ψ .

8 Estimating Roll and Pitch from Accelerometer Data

We now apply the Direction Cosine Matrix formalism to the concrete problem of estimating a device's orientation directly from raw accelerometer measurements, while the device is stationary, or more precisely, whenever the dynamic component of acceleration is negligible compared to gravity.

Let $L(X, Y, Z)$ denote a level reference frame, in which the Z axis is aligned with the local vertical, pointing away from the Earth's surface, and let $T(X', Y', Z')$ denote the tilted body frame attached to the device, obtained from L by a 3-2-1 (yaw, pitch, roll) rotation.

As established in Section 3, an accelerometer at rest reports a proper acceleration whose magnitude equals the local gravitational acceleration $g \approx 9.81 \text{ m/s}^2$, and whose direction, in the level frame, points along the positive Z axis, that is, vertically upward, opposing the pull of gravity. We write this level-frame reading as the vector $[0, 0, g]^T$.

Let $[a_x, a_y, a_z]^T$ denote the corresponding reading in the tilted body frame T . Since the yaw rotation ψ is a rotation about the vertical axis, the very axis along which gravity itself acts, yaw has no effect whatsoever on the




Figure 12: Step 1: rotating the Z axis about the X axis by the roll angle ϕ .

direction of the gravity vector as seen in the level frame, and a pure yaw rotation therefore leaves the accelerometer reading completely unchanged. We may accordingly set $\psi = 0$ without loss of generality for the purpose of this derivation, and focus only on the roll and pitch rotations. Substituting $\psi = 0$ into Equation (3) gives

$$B(\phi, \theta, 0) = \begin{pmatrix} \cos \theta & 0 & -\sin \theta \\ \sin \phi \sin \theta & \cos \phi & \sin \phi \cos \theta \\ \cos \phi \sin \theta & -\sin \phi & \cos \phi \cos \theta \end{pmatrix}. \quad (4)$$

The accelerometer reading in the tilted frame is obtained by applying this DCM to the level-frame vector $[0, 0, g]^T$:

$$\begin{pmatrix} a_x \\ a_y \\ a_z \end{pmatrix} = \begin{pmatrix} \cos \theta & 0 & -\sin \theta \\ \sin \phi \sin \theta & \cos \phi & \sin \phi \cos \theta \\ \cos \phi \sin \theta & -\sin \phi & \cos \phi \cos \theta \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ g \end{pmatrix}. \quad (5)$$

Carrying out this multiplication gives the three component equations¹

$$a_x = -g \sin \theta, \quad (6)$$

$$a_y = g \sin \phi \cos \theta, \quad (7)$$

$$a_z = g \cos \phi \cos \theta. \quad (8)$$

As a sanity check, note that when the device is perfectly level, $\phi = \theta = 0$, these equations correctly reduce to $a_x = a_y = 0$ and $a_z = g$, matching the reading obtained from a stationary, level accelerometer in Section 3.

Equations (6) to (8) may now be inverted to recover the roll and pitch angles directly from a triaxial accelerometer reading. From Equation (6) we obtain $\theta = \sin^{-1}(-a_x/g)$; however, this arcsine based formula becomes numerically ill-conditioned near $\theta = \pm 90^\circ$. A more robust equivalent expression follows on noting that $\sqrt{a_y^2 + a_z^2} = g \cos \theta$, which uses $\sin^2 \phi + \cos^2 \phi = 1$ to eliminate ϕ :

$$\theta = \tan^{-1} \left(\frac{-a_x}{\sqrt{a_y^2 + a_z^2}} \right). \quad (9)$$

¹The sign convention adopted here for a_x is chosen so as to remain consistent with the reading of a stationary, level accelerometer established in Section 3, namely $a_z = +g$ when $\phi = \theta = 0$.




Figure 13: Step 2: rotating the X axis about the Y axis by the pitch angle θ .

Similarly, dividing Equation (7) by Equation (8) eliminates θ entirely and yields the roll angle

$$\phi = \tan^{-1} \left(\frac{a_y}{a_z} \right). \quad (10)$$

It is worth emphasising precisely what these equations do, and do not, allow us to compute. Using the accelerometer alone, we can determine both the roll and the pitch angle of the device relative to the level frame, but the yaw angle remains entirely undetermined. This is a direct and unavoidable consequence of the fact, noted above, that gravity always points along the (level) Z axis: a rotation about this same axis, that is, precisely a yaw rotation, has no observable effect on the accelerometer readings, since it never moves the gravity vector out of alignment with Z. To determine the yaw angle we must therefore turn to a second sensing modality.

9 Measuring Yaw: The Role of the Gyroscope

A gyroscope is a device that measures angular velocity, the instantaneous rotation rate of the object about each of its three axes, expressed with respect to an inertial frame of reference.

Since the roll and pitch angles are already available from the accelerometer, as derived in Section 8, the gyroscope's measurements of angular velocity about the three axes can, in principle, be used to additionally recover the yaw angle, thereby completing the full three-angle orientation estimate.

Concretely, the gyroscope reports a triaxial angular velocity vector $[\omega_x, \omega_y, \omega_z]$. Multiplying each component by a small time interval δt gives the incremental angular displacement over that interval; accumulating, that is, integrating, these increments over time yields the three orientation angles: roll, pitch and yaw.

In practice, orientation estimates obtained purely by integrating gyroscope readings in this manner tend to accumulate substantial noise and drift over time, arising primarily from the finite precision with which the sampling interval δt , and the angular velocity itself, can be measured, so that small errors accumulate steadily over the course of the integration. This shortcoming motivates combining the gyroscope with the other IMU sensors, an approach we take up in the next section.

10 Sensor Fusion for Robust Orientation Estimation

We conclude this part by consolidating the complete set of relations used to estimate a device's three dimensional orientation from IMU data.




Figure 14: Step 3: rotating the Y and X axes about the Z axis by the yaw angle ψ , completing the transformation.

Given a triaxial accelerometer reading $[a_x, a_y, a_z]$, the magnitude of the locally sensed gravitational acceleration can be recovered directly as

$$g = \sqrt{a_x^2 + a_y^2 + a_z^2}, \quad (11)$$

which offers a convenient way to check that the sensor is, approximately, experiencing only the gravity offset with negligible dynamic acceleration, a condition implicitly assumed throughout the derivation of Section 8. Under this condition, the roll and pitch angles are computed exactly as derived previously:

$$\theta = \tan^{-1} \left(\frac{-a_x}{\sqrt{a_y^2 + a_z^2}} \right), \quad \phi = \tan^{-1} \left(\frac{a_y}{a_z} \right). \quad (12)$$

The yaw angle, being unobservable from the accelerometer alone, is typically obtained by combining gyroscope integration, for short-term accuracy and responsiveness, with an independent absolute reference. A magnetometer, which senses the local components $[m_x, m_y, m_z]$ of the Earth's magnetic field, provides exactly such a reference, playing a role for yaw analogous to the one gravity plays for roll and pitch. Once the raw magnetometer reading has been tilt-compensated, that is, corrected for the device's already-known roll and pitch so that m_x and m_y correctly reflect the horizontal components of the field, the yaw angle can be estimated as

$$\psi = \tan^{-1} \left(\frac{m_y}{m_x} \right). \quad (13)$$

In a full IMU implementation, the roll and pitch estimates derived from the accelerometer, together with the yaw estimate derived from the magnetometer, serve as slowly varying, drift-free reference signals. These are fused, using a Kalman filter or a complementary filter, with the fast, low-latency, but drift-prone orientation estimate obtained by continuously integrating the gyroscope's angular velocity readings. The resulting estimate combines the low noise and long-term stability of the accelerometer and magnetometer with the responsiveness of the gyroscope, and this fusion approach forms the standard basis of essentially all commercial IMU-based orientation tracking systems.




Figure 15: Two reference frames: the fixed frame R_1 of an external observer, and the local frame R_2 attached to the object.

11 Looking Ahead: Gyroscopic Effects and Self-Balancing

The discussion so far has treated the gyroscope purely as a rate sensor, used to estimate orientation via integration. Gyroscopes, however, also exhibit a range of distinctive mechanical behaviours, most notably gyroscopic precession and the associated self-stabilising properties, that have been exploited historically in a variety of ingenious engineering systems.

A classic illustration is the Brennan monorail, an early twentieth century single-rail vehicle that used large internal spinning gyroscopes to remain upright and balanced on a single rail, a feat that would otherwise be mechanically unstable.

Understanding how a spinning gyroscope can be harnessed to actively maintain balance, and more generally, how the dynamics of rotating rigid bodies can be exploited for stabilisation and control, is the subject we turn to next.




Figure 16: The level frame $L(X,Y,Z)$ and the tilted frame $T(X',Y',Z')$, related by the roll angle ϕ and the pitch angle θ , with gravity $g = 9.81 \text{ m/s}^2$ acting vertically downward.




Figure 17: A mechanical gyroscope, showing the gyroscope frame, gimbal, rotor and spin axis.