IMU Calibration and Orientation Estimation on the Wombat Educational Robot

Tobias Madlberger
HTBLuVA St. Pölten
St. Pölten, Austria
tobias.madlberger@gmail.com

Matthias Greil
HTBLuVA St. Pölten
St. Pölten, Austria
matthias.greil@gmail.com

***Abstract*—Low-cost MEMS inertial measurement units on educational robot platforms suffer from systematic biases, scale errors, and magnetic distortions that degrade orientation accuracy if left uncorrected. This paper presents an equipment-free calibration pipeline for the MPU-9250 IMU on the KIPR Wombat controller, covering accelerometer offset and scale correction as well as magnetometer hard- and soft-iron compensation via ellipsoid fitting. Four orientation estimation methods—raw gyroscope integration, the on-chip Digital Motion Processor (DMP), and a Mahony complementary filter in 6-axis and 9-axis configurations—are evaluated in static and dynamic tests. The DMP's fusion adds unexpected yaw drift beyond the gyroscope's natural bias, while the Mahony 9-axis filter achieves the best static heading stability but degrades during fast rotation. For dynamic heading, the DMP tracks best due to its higher on-chip update rate. All calibration scripts and data are openly available.**

***Index Terms*—IMU calibration, sensor fusion, digital motion processor, educational robotics, Botball, MEMS sensors**

I. Introduction

Accurate orientation is essential for autonomous navigation. In Botball, where robots must move precisely within a two-minute match, even small angular errors add up during dead-reckoning: a gyroscope bias of 0.05 °/s alone causes 3° of heading error per minute, or several centimetres of sideways offset.

The KIPR Wombat controller—the standard Botball platform—uses an MPU-9250 nine-axis MEMS IMU [1]. Low-cost sensors like this suffer from biases, scale errors, and magnetometer hard- and soft-iron distortions from nearby motors and metal parts, all of which degrade dead-reckoned heading if left uncorrected. Despite this, there is little published guidance on calibrating an IMU on platforms like the Wombat; most teams use factory defaults or accept the drift.

The MPU-9250 additionally provides a Digital Motion Processor (DMP) for on-chip 6-axis fusion and a Motion Processing Library (MPL) that extends the DMP to 9-axis with magnetometer data [2]. Both are poorly documented, and the Wombat's stock firmware enables neither. The authors wrote custom firmware that activates the DMP and exposes the magnetometer, enabling a direct comparison of hardware and software fusion.

This paper makes the following contributions:

- A calibration method for the MPU-9250's accelerometer and magnetometer that works without special equipment.
- A comparison of four orientation estimation methods, showing that the DMP's own fusion adds more yaw drift than the gyroscope's natural bias alone.
- Evidence that a 9-axis Mahony filter with magnetometer gives the best static heading stability, converging to magnetic north within 1.3 seconds and drifting less than 1 °/h.

II. State of the Art

A. IMU Error Modelling and Calibration

The systematic errors of MEMS inertial sensors are well characterised. The standard model, used by Frosio et al. [3] and Tedaldi et al. [4], expresses the measured output as $\boldsymbol{y} = \boldsymbol{T}_m \boldsymbol{K}(\boldsymbol{x} + \boldsymbol{b}) + \boldsymbol{n}$, where $\boldsymbol{x}$ is the true physical quantity, $\boldsymbol{b}$ the bias vector, $\boldsymbol{K}$ the scale factor matrix, $\boldsymbol{T}_m$ the axis misalignment, and $\boldsymbol{n}$ noise. Full identification requires carefully designed motion sequences [5]. For low-cost IMUs, misalignment is typically below 0.5° and may be neglected [2], reducing the problem to bias and scale.

Magnetometer calibration requires additional treatment because ferromagnetic components in the robot body introduce hard-iron (constant offset) and soft-iron (orientation-dependent distortion) errors. Ellipsoid fitting maps distorted measurement data back to a unit sphere and corrects both effects simultaneously [6]. Outlier rejection based on the median absolute deviation (MAD) is applied prior to fitting to suppress transient magnetic anomalies during data collection.

B. Sensor Fusion for Orientation Estimation

Complementary filters combine sensors with opposite strengths: gyroscopes are accurate over short timescales but drift, while accelerometers and magnetometers are noisy short-term but stay honest over the long run. Mahony et al. [7] turned this idea into a concrete 3D orientation filter: each step, it advances the orientation estimate with the gyroscope and then nudges it back toward the direction that gravity (and optionally magnetic north) point to, while simultaneously learning and subtracting the gyroscope's slowly varying bias—all with only two tuning gains and mathematically proven stability. Madgwick [8] later proposed a gradient-descent variant with a single tuning parameter. The Extended Kalman Filter [9, 10] offers a probabilistic alternative that can incorporate full noise-covariance information, at higher computational cost.

For resource-constrained embedded systems typical of educational robotics, the computational efficiency and stability of gradient-free complementary filters is advantageous.

C. Hardware-Integrated Fusion

The MPU-9250's Digital Motion Processor (DMP) runs a proprietary 6-axis fusion algorithm on the sensor chip, combining gyroscope and accelerometer data without using the host CPU. InvenSense's Motion Processing Library (MPL) adds magnetometer data for 9-axis orientation. Both are poorly documented—InvenSense provides driver code but does not describe the algorithms or their limitations.

D. Gap in the Literature

Prior ECER work has explored inertial sensing for Botball: Klauninger et al. [11] demonstrated that gyroscope-based correction improves navigation accuracy on the Wombat, but applied only runtime offset subtraction without systematic calibration of scale errors, magnetometer distortions, or noise characterisation. Earlier GCER papers on the Wallaby platform—Torbati and Zhang [12] on manual bias subtraction, Hu [13] on elapsed-time integration, and Guzmits et al. [14] on accelerometer-based distance measurement—similarly used raw sensor data without calibration or fusion. No prior work in the ECER or Botball community has addressed a complete IMU calibration pipeline or structured filter selection for the Wombat platform. This paper addresses that gap.

III. Concept and Design

A. Accelerometer Calibration

Accelerometer errors follow the standard model from Section II. The dominant effects are a *bias* (a constant offset at rest) and a *scale error* (axes that over- or under-report motion, e.g., 1 g of gravity reading as 2 g). For the Wombat, axis misalignment is neglected ($< 0.5^\circ$ per datasheet [2]), reducing the model to a bias vector $\boldsymbol{b}_a \in \mathbb{R}^3$ and a diagonal scale matrix $\boldsymbol{K}_a$:

$$\boldsymbol{a}_{\text{corr}} = \boldsymbol{K}_a^{-1}(\boldsymbol{a}_{\text{raw}} - \boldsymbol{b}_a) \tag{1}$$

Biases are found by averaging static measurements in a known orientation; scale factors come from the ratio of measured to expected gravity. Validation is geometric: because gravity has constant magnitude, a calibrated accelerometer rotated through all orientations should trace a 1 g sphere. Ellipsoid fitting on free-rotation data measures how close the readings come to that ideal [4].

B. Magnetometer Calibration

Magnetometers suffer two characteristic distortions. *Hard-iron* effects—permanent magnetic sources rigidly attached to the robot (e.g., magnetised screws, motor magnets at rest)—add a constant offset to every reading. *Soft-iron* effects come from ferromagnetic material that distorts the ambient field depending on orientation, stretching and tilting the apparent field. The combined error model uses a hard-iron offset $\boldsymbol{b}_m$ and a soft-iron matrix $\boldsymbol{S}_m$:

$$\boldsymbol{m}_{\text{raw}} = \boldsymbol{S}_m \boldsymbol{m}_{\text{true}} + \boldsymbol{b}_m + \boldsymbol{n}_m \tag{2}$$

Ellipsoid fitting exploits a simple fact: rotated through all orientations, an ideal magnetometer traces a sphere with radius equal to the local field magnitude. Hard-iron offsets shift that sphere off the origin; soft-iron effects squash it into a tilted ellipsoid. Fitting the best ellipsoid therefore recovers both corrections at once. Least-squares fitting runs on data from full 3D rotation, with MAD-based outlier rejection for transient anomalies. At runtime:

$$\boldsymbol{m}_{\text{corr}} = \boldsymbol{S}_m^{-1}(\boldsymbol{m}_{\text{raw}} - \boldsymbol{b}_m) \tag{3}$$

This calibration must be repeated when the robot's setup changes, because motor currents and ferromagnetic parts change both $\boldsymbol{b}_m$ and $\boldsymbol{S}_m$ [15].

C. Orientation Estimation Approaches

The authors evaluate three approaches in four configurations:

Digital Motion Processor (DMP): The MPU-9250's built-in DMP runs a proprietary fusion algorithm on the sensor chip, combining gyroscope and accelerometer data at the hardware level (6-axis). It does its own gyroscope bias calibration internally.

InvenSense MPL: The Motion Processing Library is designed to extend the DMP by additionally fusing magnetometer data (9-axis). In the authors' configuration, the MPL's compass integration did not function correctly, so its output is effectively 6-axis and serves as a second hardware-fusion data point.

Mahony filter: A software sensor-fusion filter [7] that runs on the Raspberry Pi. The idea is intuitive: integrating the gyroscope alone gives a smooth short-term orientation but slowly drifts away from the truth, while gravity (from the accelerometer) and magnetic north (from the magnetometer) are noisy but do not drift. Each step, the filter rotates the previous orientation forward with the gyroscope and then nudges the result back toward the direction where gravity and magnetic north agree. The size of that nudge is set by a PI (proportional–integral) controller acting on the orientation error: the proportional gain k_P reacts to the instantaneous error and pulls the estimate toward the reference, while the integral gain k_I accumulates the error over time and thereby learns the gyroscope's slowly varying bias and subtracts it online—so no separate gyroscope bias calibration is needed. The authors chose this filter because it is simple, well-established, and requires only two scalar gains instead of full noise covariance tuning. It is tested in 6-axis (gyroscope + accelerometer) and 9-axis (additionally magnetometer) configurations on the calibrated data.

Raw gyroscope integration: Direct numerical integration of angular rates, serving as a baseline.

Note that the DMP performs its own internal calibration, while the accelerometer and magnetometer calibrations described above are applied only to the Mahony filter inputs.

IV. Implementation

A. Hardware Platform

The Wombat pairs a Raspberry Pi 3 (user programs) with an STM32 microcontroller (real-time hardware control) [1]. The MPU-9250 connects to the STM32 over SPI; the gyroscope runs at ±2000 °/s with a 42 Hz digital low-pass filter. The DMP samples at 200 Hz internally, the STM32 drains the FIFO at 50 Hz, and the data reaches the Pi at about 75 Hz through a shared SPI buffer.

B. Calibration Data Collection

Accelerometer calibration used 30,000+ static samples (z vertical) for bias and scale plus 25,439 free-rotation samples for ellipsoid validation. Magnetometer calibration used 12,805 samples from eight minutes of full 3D rotation. Gyroscope bias is estimated at start-up by the DMP's built-in calibration.

C. Calibration Implementation

Calibration parameters are computed in Python. Magnetometer ellipsoid fitting runs a least-squares solver on the general quadric form, applied to raw AK8963 ADC output (including the factory sensitivity adjustment registers); the resulting soft-iron matrix absorbs both physical distortion and internal unit scaling.

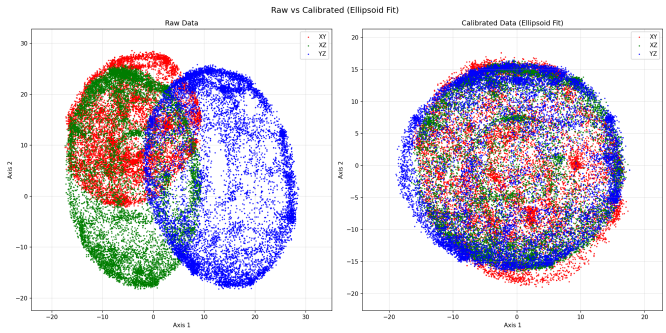

Fig. 1. Magnetometer data before (left) and after (right) ellipsoid calibration. Raw data forms a shifted, distorted ellipsoid; calibrated data approximates a centred sphere. 12,805 measurements per projection.

D. DMP and MPL Configuration

The DMP is configured for 6-axis low-power quaternion estimation (gyroscope + accelerometer) with built-in gyroscope calibration. Its quaternion is read from the FIFO and forwarded to the Pi over SPI.

InvenSense's MPL is also enabled on the STM32 to fuse the DMP's output with magnetometer data into a 9-axis quaternion. In the authors' configuration it tracked the DMP within 0.01° throughout all tests, confirming it ran as a 6-axis filter. Both quaternions are forwarded to the Pi.

E. Mahony Filter Implementation

The Mahony filter runs offline in Python. Its gyroscope and accelerometer inputs come from the firmware already preprocessed by the DMP (bias-corrected gyro, scaled and gravity-compensated accel)—cleaner than truly raw data, which partly explains its good performance. Gains are $k_P = 2.0$, $k_I = 0.005$ from manual tuning. Two configurations are tested: 6-axis (gyro + accel) and 9-axis (adding magnetometer).

The magnetometer requires axis realignment: the AK8963 has its X and Y axes swapped relative to the gyroscope/accelerometer frame [2].

V. Results and Discussion

A. Calibration Parameters

Accelerometer: Static z-vertical measurements gave a bias $\boldsymbol{b}_a = [0.236, -0.099, 0.000]^T$ m/s² and a z-axis scale factor $k_z = 1.981$—the raw z reading was 19.43 m/s², almost twice the expected 9.81. After correction, ellipsoid fitting on 25,439 free-rotation samples confirms the data form a unit sphere centred at the origin (Fig. 2).

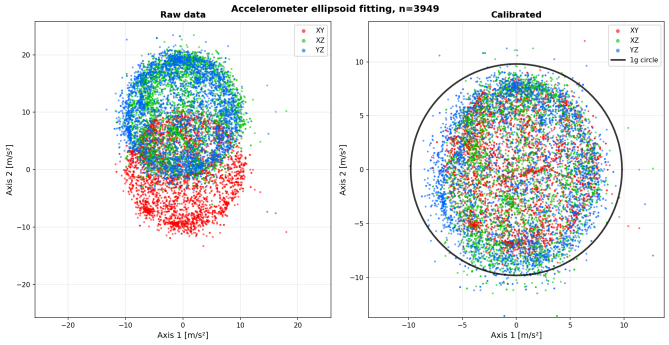

Fig. 2. Accelerometer data before (left) and after (right) ellipsoid calibration. Raw data forms a shifted, elongated ellipsoid due to the z-axis scale error; calibrated data approximates a sphere at the expected 1g radius. 25,439 measurements per projection.

Magnetometer: Ellipsoid fitting to 12,805 samples (after outlier rejection) yielded a hard-iron offset $\boldsymbol{b}_m = [-3.86, -4.38, 14.87]^T \sim \mu$T—the z-component alone is about 30% of the local geomagnetic field (48 µT in Central Europe)—and a soft-iron diagonal $[0.067, 0.063, 0.045]$ (Fig. 1).

B. Static Orientation Comparison

The Wombat was kept still for 10 minutes while recording raw sensor data and the DMP quaternion at about 75 Hz. The Mahony filter and raw gyroscope integration were computed offline from body-frame data reconstructed by inverse-rotating the DMP's world-frame output—already bias-corrected by the DMP, so the comparison reflects the actual pipeline rather than truly raw sensor data. MPL 9-axis fusion is excluded: despite enabling all required features, the precompiled InvenSense library did not produce 9-axis output in the authors' firmware configuration. Table I and Fig. 3 show the results.

TABLE I

ORIENTATION CHANGE AFTER 10 MINUTES OF STATIC RECORDING. THE MAHONY (9AX) "YAW" COLUMN SHOWS INITIAL CONVERGENCE TO MAGNETIC NORTH, NOT DRIFT—SEE TEXT.

Method	**Roll** [°]	**Pitch** [°]	**Yaw** [°]	**Yaw drift** [°/h]†
Mahony (9ax)	−0.16	−1.52	−6.87	**0.19**
Mahony (6ax)	−0.11	−1.51	3.17	21.1
Raw gyro	−0.54	10.12	2.98	26.8
DMP (6ax)	0.05	0.31	−19.45	109.5

†Yaw drift rate computed from the last 8 minutes (after the Mahony 9-axis filter has converged).

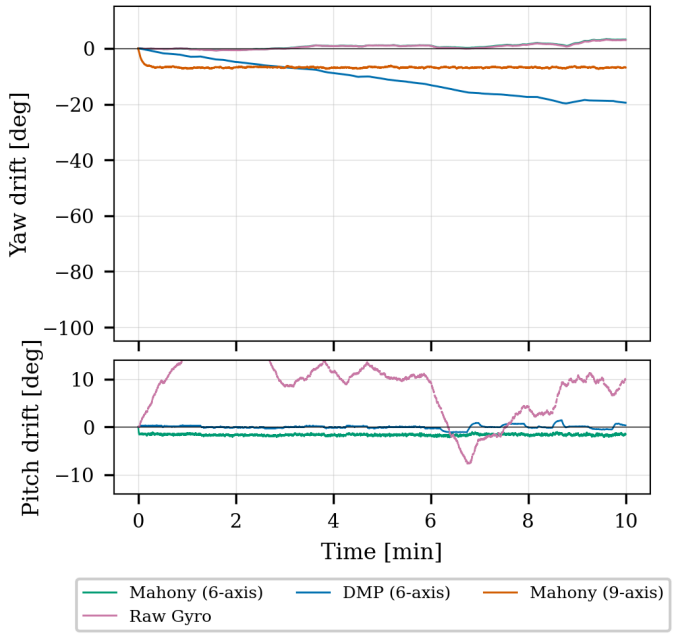

Fig. 3. Yaw drift (top) and pitch drift (bottom) over 10 minutes. The Mahony 9-axis curve converges to magnetic north in the first 1.3 seconds; after settling, it is the most stable method. The DMP drifts steadily at −2 °/min. All fusion methods control pitch; raw integration does not.

Roll and pitch are kept stable by all fusion methods. The DMP has the best pitch result (0.31°), while all Mahony variants stay below 1.6°. Raw gyroscope integration drifts 10° in pitch, which shows that gravity-based correction is needed.

Yaw drift is the main source of error for the 6-axis methods. The Mahony 6-axis filter and raw gyroscope integration both drift about 3° in yaw over 10 minutes, while the DMP drifts 19°.

Mahony 9-axis: convergence to magnetic north. The 6.9° yaw change is not drift—it is the filter converging from its initial identity quaternion to the magnetic-north heading in about 1.3 seconds. After that, yaw changes by only 0.08° over the remaining 8 minutes (0.19 °/h)—the most stable method by a large margin, about 150 times better than the DMP. The magnetometer data here was uncalibrated (raw AK8963 ADC values); the Mahony filter's internal normalization makes it robust to bias, since heading comes from the field *direction*, not magnitude. Motors were off throughout, so motor-induced distortion was not a factor in this static test.

C. Why the DMP Drifts More Than Mahony

The DMP's 19° yaw drift is surprising given that raw gyroscope integration drifts only 3°. Both raw integration and the Mahony filter use DMP-preprocessed (bias-corrected) gyroscope data, whose mean z-axis bias of 0.005 °/s matches the observed 3° drift over 10 minutes exactly. The DMP's extra 16° must therefore come from its own fusion.

The most likely cause is gravity-vector coupling: when the DMP corrects roll and pitch from the accelerometer, small errors in the gravity estimate leak into yaw—a known pitfall of 6-axis fusion when the gravity vector is not perfectly axis-aligned. The Mahony filter, with gentler PI gains, avoids this. Fig. 3 confirms the pattern: the DMP drifts steadily at −2 °/min while the Mahony filter and raw gyroscope stay near zero.

D. Dynamic Heading Validation

A smartphone (Samsung Galaxy Z Flip3) was mounted firmly on the robot, and both devices recorded for 2 minutes of yaw-only rotation on a flat surface similar to typical Botball match conditions. The smartphone's IMU introduces an unquantified error floor in the reference heading. Time alignment used cross-correlation of gyroscope rate magnitudes ($r = 0.95$), confirming that both devices were rigidly coupled. Results: Table II and Fig. 4.

TABLE II

DYNAMIC HEADING ACCURACY OVER 2 MINUTES OF YAW ROTATION AGAINST A SMARTPHONE REFERENCE. CORRELATION r MEASURES SHAPE AGREEMENT; MAE AND RMSE MEASURE ABSOLUTE HEADING ERROR.

Method	r	**MAE** [°]	**RMSE** [°]
DMP (6ax)	**0.998**	**12.0**	**31.9**
Raw gyro	0.998	29.8	49.7
Mahony (6ax)	0.997	31.9	51.1
Mahony (9ax)	0.989	112.9	127.0

The DMP follows the smartphone reference closely ($r = 0.998$, MAE = 12°). Mahony 6-axis and raw gyroscope integration also follow the shape well ($r > 0.99$) but accumulate about 30° of drift over 2 minutes.

Magnetometer disturbs heading during motion. The Mahony 9-axis filter has high correlation ($r = 0.989$) but the worst MAE (113°). Fig. 4 shows why: during fast rotations the magnetometer correction continuously pulls the estimate toward magnetic north, fighting the gyroscope. The high correlation shows the filter still tracks the *shape* of the motion correctly, but the absolute heading lags and overshoots. This is the opposite of the static result: the same correction that gives the lowest static drift (0.19 °/h) degrades accuracy during the rapid yaw changes typical of a Botball match.

The DMP's advantage during motion comes from running on the sensor chip at the full internal rate (200 Hz), while the Mahony filter and raw integration run offline on the 75 Hz data that arrives at the Pi.

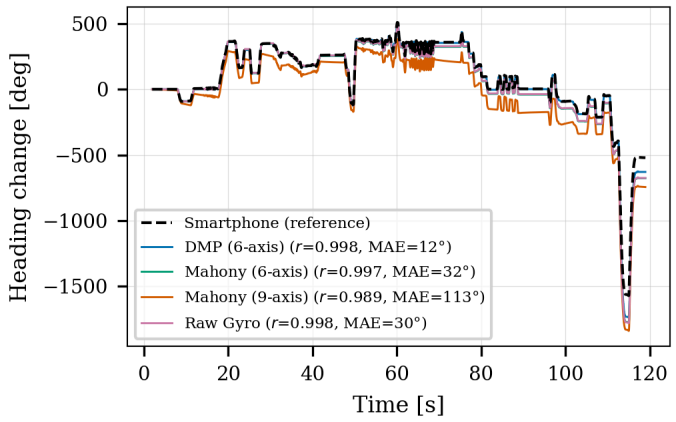

Fig. 4. Dynamic heading validation against a smartphone reference over 2 minutes of yaw rotation. The DMP follows the reference closely (MAE = 12°); Mahony 6-axis tracks the shape but accumulates drift.

VI. Conclusion

This paper compared four orientation estimation methods for the MPU-9250 IMU on the KIPR Wombat educational robot, together with accelerometer and magnetometer calibration.

A 10-minute static test showed three main results. First, all fusion methods keep roll and pitch drift below 1.6° by using gravity as a reference (raw integration: 12°). Second, the DMP adds 16° of yaw drift on top of the gyroscope's 3° natural bias—probably from accelerometer correction leaking into yaw; the Mahony 6-axis filter avoids this (3° drift). Third, adding the magnetometer (Mahony 9-axis) gives the lowest static drift (0.19 °/h), 150 times better than the DMP.

A dynamic test against a smartphone reference reversed the picture. The DMP tracks heading best during motion ($r = 0.998$, MAE = 12°), running on-chip at the full sensor rate. The Mahony 6-axis filter matches the shape ($r > 0.99$) but accumulates about 30° of drift over 2 minutes. The Mahony 9-axis filter is worst (MAE = 113°): the same magnetometer correction that stabilizes static drift fights the gyroscope during fast rotations.

A. Limitations

All measurements come from a single Wombat unit. Differences in IMU bias and magnetometer distortion between units may mean that the exact drift values do not generalize. The effect of motor-induced magnetic interference during operation was not tested and remains an open question—drive motors and servos produce magnetic fields that could degrade the 9-axis filter's heading correction during a match.

B. Future Work

Motor-induced magnetic interference is the most pressing open question: characterising how drive currents affect the magnetometer would show whether the 9-axis Mahony filter is usable during a match or only while stationary. Testing additional Wombat units would quantify hardware-to-hardware variation. Based on these results, the authors recommend the DMP for dynamic heading during motion and the Mahony 9-axis filter for stable absolute heading while stationary. All calibration scripts, data, and parameters are openly available[1] and directly reusable by other Botball teams.

References

[1] KIPR, "Wombat Educational Robotics Controller." Accessed: Apr. 10, 2026. [Online]. Available: https://www.kipr.org/kipr/hardware

[2] InvenSense Inc., "MPU-9250 Product Specification Revision 1.1." 2016.

[3] I. Frosio, F. Pedersini, and N. A. Borghese, "Autocalibration of MEMS Accelerometers," *IEEE Transactions on Instrumentation and Measurement*, vol. 58, no. 6, pp. 2034–2041, 2009, Accessed: Apr. 10, 2026. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/4655611/

[4] D. Tedaldi, A. Pretto, and E. Menegatti, "A Robust and Easy to Implement Method for IMU Calibration without External Equipments," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 3042–3049. doi: 10.1109/ICRA.2014.6907297.

[5] F. Ferraris, U. Grimaldi, and M. Parvis, "Procedure for Effortless In-Field Calibration of Three-Axis Rate Gyros and Accelerometers," *Sensors and Materials*, vol. 7, no. 5, pp. 311–330, 1995.

[6] J. F. Vasconcelos, G. Elkaim, C. Silvestre, P. Oliveira, and B. Cardeira, "A Geometric Approach to Strapdown Magnetometer Calibration in Sensor Frame," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 47, no. 2, pp. 1293–1306, 2011, doi: 10.1109/TAES.2011.5751259.

[7] R. Mahony, T. Hamel, and J.-M. Pflimlin, "Nonlinear Complementary Filters on the Special Orthogonal Group," *IEEE Transactions on Automatic Control*, vol. 53, no. 5, pp. 1203–1218, 2008, doi: 10.1109/TAC.2008.923738.

[8] S. O. H. Madgwick, "An Efficient Orientation Filter for Inertial and Inertial/Magnetic Sensor Arrays," technical report, 2010. Accessed: Apr. 10, 2026. [Online]. Available: https://web.enib.fr/~kerhoas/iot/reseau-de-capteurs/carte-imu-mpu9250/documents/INVENSENSE/madgwick_internal_report.pdf

[9] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960, doi: 10.1115/1.3662552.

[10] J. L. Crassidis, F. L. Markley, and Y. Cheng, "Survey of Nonlinear Attitude Estimation Methods," *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 1, pp. 12–28, 2007, doi: 10.2514/1.22452.

[11] B. Klauninger, K. Lindorfer, S. Kawicher, T. Koch, V. Griesmayer, and L. Kornthaler, "Enhancement of Accuracy in Botball Navigation," in *European Conference on Educational Robotics (ECER)*, Wiener Neustadt, Austria, 2023. Accessed: Apr. 10, 2026. [Online]. Available: https://ecer.pria.at/archive/ecer-2023/papers/Enhancement_of_Accuracy_in_Botball_Navigation.pdf

[12] R. Torbati and A. Zhang, "Using the Gyroscope in Botball for Simple and Consistent Navigation," in *Global Conference on Educational Robotics (GCER)*, 2018. Accessed: Apr. 10, 2026. [Online]. Available: http://files.kipr.org/gcer/2018/Using%20the%20Gyroscope%20in%20Botball.pdf

[13] Y. Hu, "Utilizing the Internal Gyroscope for Reliable Inertial Navigation," in *Global Conference on Educational Robotics (GCER)*, 2019. Accessed: Apr. 10, 2026. [Online]. Available: https://files.kipr.org/gcer/2019/Yusen%20Hu%20-%20Gyroscope.pdf

[14] M. Guzmits, C. Schnabl, M. Eiwen, S. Marton-Lindenthal, S. Rohrer, and A. Lampalzer, "Distance Measurement using a Surface Micromachined Accelerometer," in *Global Conference on Educational Robotics (GCER)*, 2019. Accessed: Apr. 10, 2026. [Online]. Available: https://files.kipr.org/gcer/2019/2019%20GCER%20Paper%20Wiener%20Neustadt.pdf

[15] M. J. Caruso, "Applications of Magnetoresistive Sensors in Navigation Systems." 1997.

[1] https://github.com/ToberoCat/wombat-imu-calibration