Applied Artificial Intelligence and Computing Systems
https://www.scigate.org
Science Gate Academic
Vol. 5, No. 6, 2026

*Article*

# Phase Prediction after Sensor Rotation with Three Reference Vectors

**Tessa Orlin*, Dario Vesk, Liora Brenn**

Department of Electrical and Computer Engineering, University of Calgary, Calgary, Canada

*Corresponding author: tessaorlin365@gmail.com

***Abstract:*** A phase predictor trained in one coordinate frame can fail when a sensor is mounted at a different orientation. We isolate this effect with an executed periodic three-axis signal experiment. A ridge model fitted to 3,000 nominal samples predicts sine and cosine phase components. Forty test trials apply independent proper rotations, each with 600 samples. Raw rotated axes, a rotation-invariant norm polynomial, and axes corrected using three noisy known reference vectors are compared. Mean phase errors are 96.62, 90.03, and 1.83 degrees; corresponding synthetic half-cycle classification accuracies are 50.07%, 50.13%, and 99.19%. A frame diagram, phase-colored three-dimensional trajectories, circular phase links, and trial-level error distributions connect calibration to prediction. The norm baseline removes orientation but also loses directional phase information. The study uses generated vectors and privileged reference poses, not human recordings, gait measurements, or a validated wearable device. It demonstrates a coordinate-calibration mechanism and its information assumptions rather than a clinical or athletic performance claim.

**Keywords:** inertial sensing; orientation calibration; periodic phase; orthogonal Procrustes; coordinate shift

## 1 A signal arrives in the wrong frame

A learning model sees numbers in a coordinate system. If the same physical motion is expressed in a different sensor frame, its axis values change even though the motion has not. A predictor trained only in the original frame may interpret that coordinate change as a different phase or activity. This paper examines the issue at the level of a generated periodic vector, where the rotation is known to the evaluator and can be estimated from a small reference collection. The aim is to identify what calibration contributes before introducing a complex temporal architecture.

The motivating setting is wearable or sports sensing, but the experiment contains no person, body segment, inertial measurement unit, or recorded movement. Its three coordinates are an illustrative signal with a periodic structure. The classification target is a synthetic half-cycle indicator defined by the sign of sine phase. It is sometimes useful as a contact-like proxy, but it is not measured foot contact, stance, or a clinical gait label. We retain precise language so that a numerical accuracy cannot acquire an unsupported physiological interpretation.

Three representations are compared. Raw axes preserve all directional information but use the wrong frame after rotation. A polynomial of vector norm removes frame orientation but also removes part of the phase information. Calibrated axes use an estimated proper rotation to return observations toward the training frame, preserving directional content. These choices expose a tradeoff between invariance and information. A representation can be insensitive to a nuisance transformation while becoming insufficient for the target, so invariance alone is not an adequate objective.

## 2 Orientation and activity learning in context

Complementary orientation filters [1], gradient-based IMU orientation [2], and quaternion filtering [3] provide established methods for inertial orientation estimation. Joint-angle measurement [4], wearable motion analysis [5], and sensor-to-segment calibration reviews [6] motivate the practical importance of reference frames. Our experiment does not implement a time-recursive orientation filter, integrate angular velocity, or estimate gravity during motion. It estimates one fixed rotation per trial from known vectors. That narrower task allows an exact coordinate argument without implying a complete inertial navigation solution.

Deep convolutional and recurrent activity recognition [7], learning surveys [8], and user-annotated acceleration recognition [9] show the broader range of wearable prediction tasks. The present ridge predictor is not intended as a competitive activity-recognition architecture. Kabsch rotation estimation [10] and least-squares transformation methods [11] supply the geometric background for aligning reference vectors. Wearable rehabilitation systems [12], activity-monitoring reviews [13], inertial smoothing [14], and efficient deep recognition [15] provide application context, while random forests [16] and general pattern recognition [17] are alternative model references rather than executed baselines.

A 2026 attitude-aided accelerometer-calibration study [18] illustrates ongoing interest in using orientation information for efficient calibration. We do not reproduce its sensor-error model or hardware experiment. Our reference vectors contain independent noise but no bias, scale-factor error, nonorthogonality, temporal offset, or mounting drift. They are privileged known canonical directions. The comparison therefore measures what a small amount of correct frame information can accomplish in a favorable model, not how easily such information can be acquired in the field.

## 3 The periodic vector and the nominal predictor

The canonical signal has a sine first coordinate with amplitude 1.1, a cosine second coordinate with amplitude 0.7, and a third coordinate equal to 0.3 times the sine of double phase plus 0.4. This creates a nonplanar periodic loop with directional structure. Training phases are independent uniform draws from negative pi to pi. Independent coordinate noise has standard deviation 0.03. There are 3,000 nominal samples, and seed 4001 fixes the experiment. All signal values are arbitrary units rather than acceleration in physical units.

$$a(\phi) = [\, 1.1\sin\phi, 0.7\cos\phi, 0.3\sin(2\phi) + 0.4 \,]^T \tag{1}$$

The nominal model appends an intercept to the noisy three-axis vector and uses ridge regression with penalty 0.001 to predict sine and cosine phase components. In a noise-free canonical frame, the first two axes already contain a linear representation of those components. The model class is therefore adequate for the intended task before rotation. Phase is recovered from the predicted components using the two-argument arctangent, and angular error is wrapped to the interval from negative pi to pi before taking its absolute value.

$$\hat{\phi} = atan2(\hat{s}, \hat{c}) \tag{2}$$

The norm baseline uses powers zero through four of the vector norm and ridge penalty 0.01 to predict the same two components. It is fitted on the same nominal training samples. Norm is unchanged by a proper rotation in the absence of noise, which gives this baseline a strong invariance property. However, the norm of the selected loop does not uniquely encode signed phase. A low-dimensional invariant can discard exactly the directional distinctions needed by the target. This baseline is included to expose that loss rather than to represent the strongest invariant learning method.

## 4 A fixed rotation and three known directions

Each of forty trials draws a proper three-dimensional rotation from a corrected Gaussian-QR construction. The diagonal signs from the triangular factor are applied before enforcing positive determinant, avoiding the directional bias that an uncorrected QR convention can introduce. Each test trial contains six hundred evenly spaced phases over three cycles. The canonical vector is rotated and independent coordinate noise with standard deviation 0.03 is added. The rotation is constant within a trial and independent across trials; time-varying sensor motion relative to the frame is not modeled.

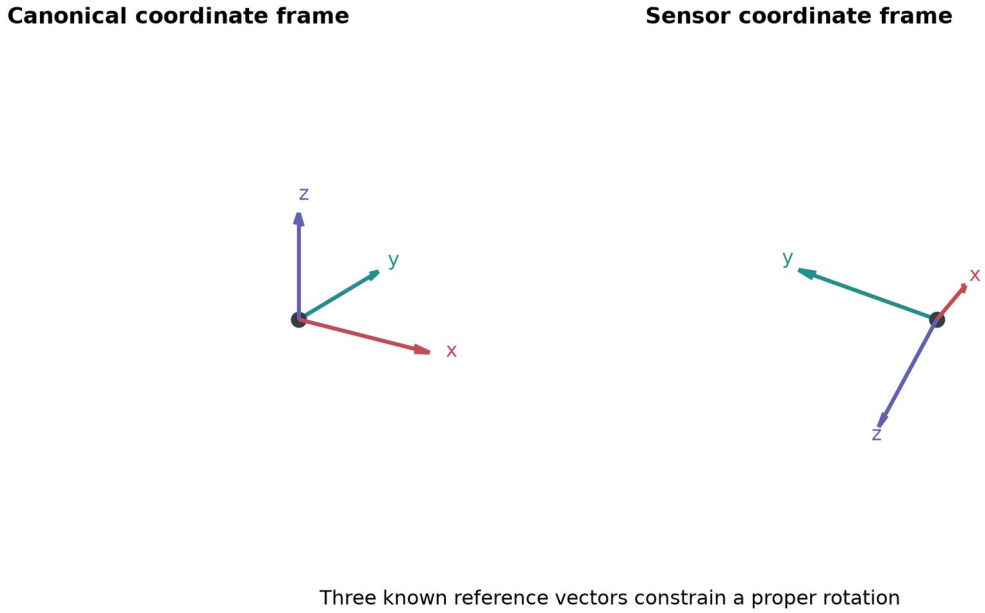



Figure 1. Canonical and sensor coordinate triads for the first generated trial. Axis colors preserve correspondence across frames. The diagram is a coordinate illustration, not an anatomical sensor-placement schematic.

The calibration set consists of three canonical unit reference vectors. Their observed versions are rotated by the same trial rotation and perturbed with independent noise of standard deviation 0.015. An orthogonal Procrustes calculation estimates a proper rotation from these three correspondences. Singular-value decomposition supplies the solution with a determinant correction that excludes reflection. This uses labeled vector identities: the algorithm knows which observed direction corresponds to which canonical direction. Without those labels, a different correspondence problem would remain.

$$C = U\,B\,V^{T},\ \hat{R} = U\,diag\left(1, 1, det\left(U\,V^{T}\right)\right)V^{T} \tag{3}$$

Here C is the cross-covariance formed from the observed and canonical reference matrices. Under the row-vector convention used in the implementation, a rotated observation is multiplied on the right by the estimated rotation to return it toward the canonical frame. The same nominal ridge coefficients are then applied. No phase labels from the test trial are used to refit the predictor. The improvement therefore comes from a changed input coordinate frame, while the predictive mapping itself remains fixed.

The reference set is small but information-rich. Three approximately orthogonal labeled vectors constrain all rotation axes in this model. They are not equivalent to three arbitrary samples from an unknown motion trajectory, whose canonical directions might be unavailable or nearly collinear. The calibration method's favorable result therefore rests on access to known reference poses. A practical wearable protocol would need to establish those poses reliably and evaluate errors in their execution, rather than assuming that the mathematical unit vectors can be collected directly.

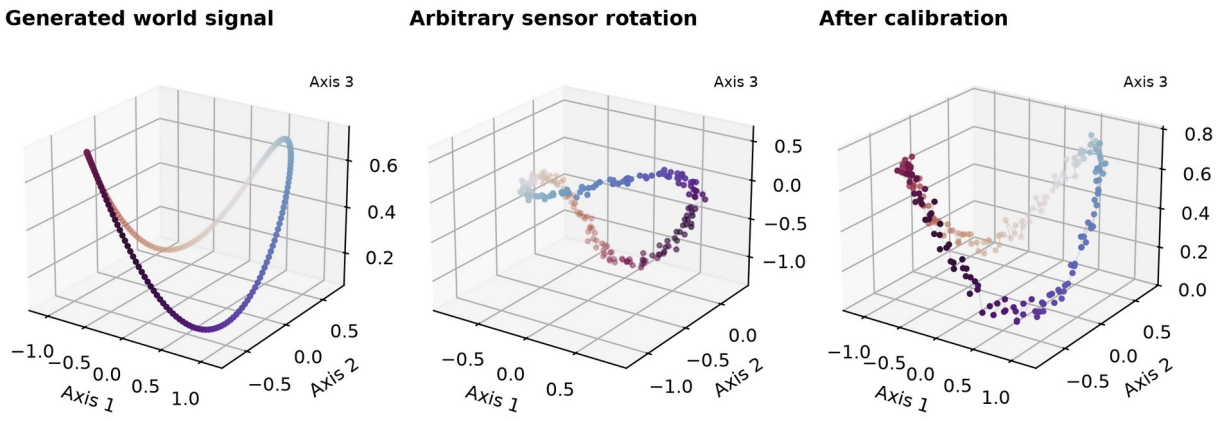



Figure 2. Phase-colored three-dimensional trajectories for the canonical signal, observed rotated signal, and calibrated signal in the first trial. Only the first cycle is displayed. Noise creates a narrow cloud around each loop; color encodes generated phase rather than a measured biomechanical quantity.

## 5 Circular prediction and the observed errors

Mean angular error is computed per trial and then averaged across forty rotations. Component RMSE compares predicted sine-cosine coordinates with their true values. The half-cycle score compares the sign of the predicted sine component with the true sign. These metrics answer different questions. A phase prediction can be numerically close while crossing a sign boundary near zero, and a half-cycle label can be correct while phase remains far from its true value. Reporting both avoids letting a coarse label conceal poor continuous timing.

Table 1. Mean performance across 40 generated rotations, 600 samples per trial.

| Representation | Phase MAE | Half-cycle accuracy | Component RMSE |
|---|---|---|---|
| Raw axes | 96.6190° | 50.07% | 1.03648 |
| Norm polynomial | 90.0303° | 50.13% | 0.70729 |
| Calibrated axes | 1.8326° | 99.19% | 0.03866 |

The raw-axis result is near chance on the coarse half-cycle task because the nominal coordinate interpretation is generally inappropriate after an arbitrary rotation. The norm baseline is also near chance, despite being rotation-invariant. Its component error near 0.707 is consistent with predicting little signed phase information in this symmetric sampling setting. Calibrated axes reduce mean angular error to approximately 1.83 degrees and retain high half-cycle accuracy. This is evidence for restoring a useful coordinate representation under the specified reference assumptions, not for a universal sensor-calibration guarantee.

The mean rotation-estimation error is approximately 1.05 degrees across trials. The phase error is somewhat larger because reference noise is only one source of prediction error: test-coordinate noise, nominal model fitting, and the loop's geometry also contribute. These two angular measures should not be subtracted as if independent scalar errors. A small frame error can affect different phases differently depending on the signal direction and the learned coefficients. Figure 4 combines trial-level distributions with one within-cycle error trace to show that distinction.

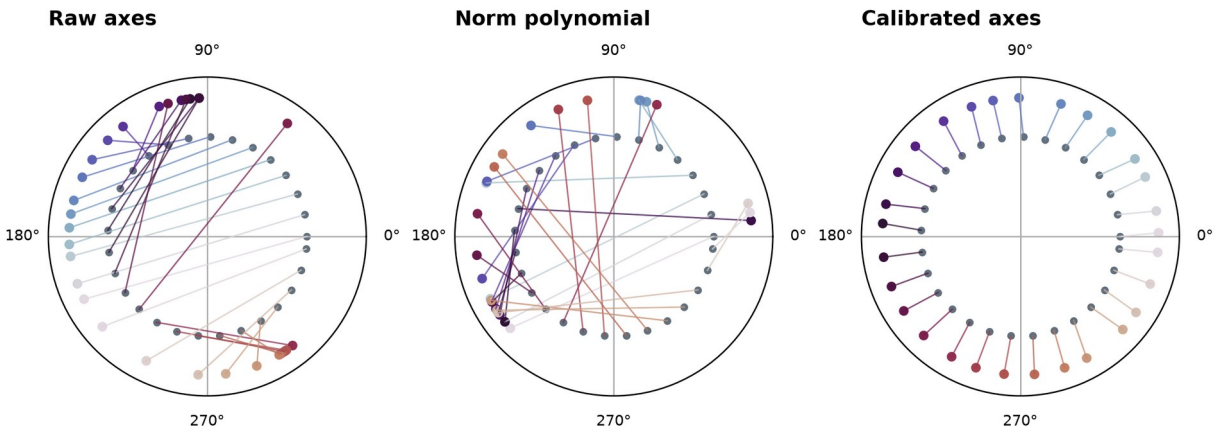



Figure 3. Circular links from true phase on the inner ring to estimated phase on the outer ring for selected first-cycle samples. Panels compare raw axes, norm polynomial, and calibrated axes in the first trial. Link length is illustrative; quantitative angular errors use wrapped differences.

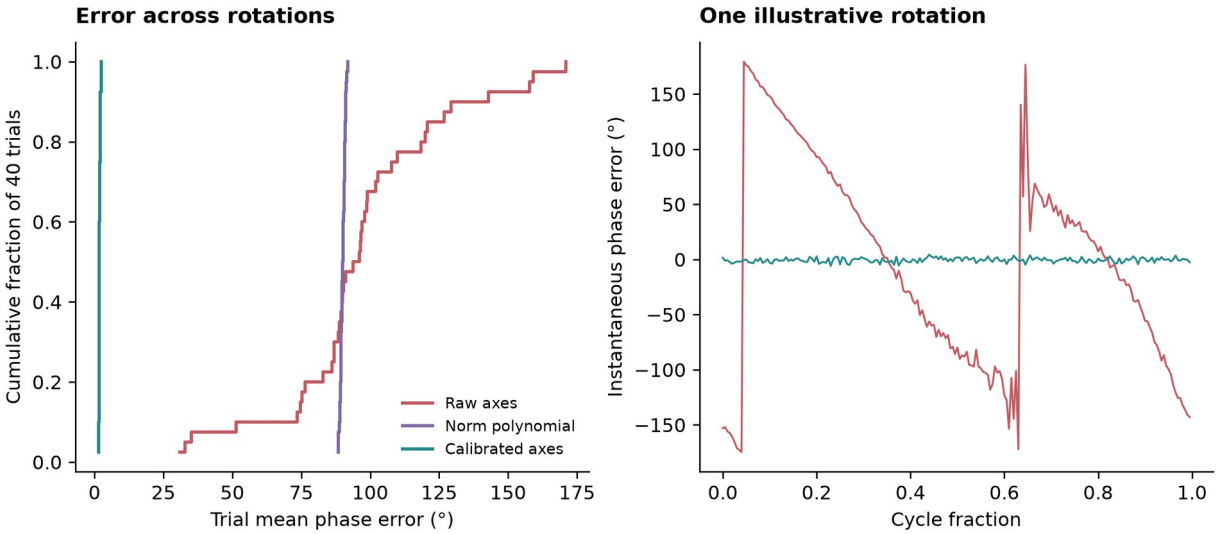



Figure 4. Empirical cumulative distribution of trial mean phase errors and an instantaneous wrapped-error trace for the first trial. Forty independent rotations determine the left panel; the six hundred samples within a trial are correlated through the same rotation and calibration estimate.

## 6 The information lost by a norm

A norm treats all vectors on a sphere as equivalent. That is exactly why it removes rotation, but it also removes direction. In our signal, multiple phases can have the same norm while requiring different sine and cosine outputs. No deterministic function of that norm can assign all such phases correctly without additional information. Increasing the polynomial degree could fit noise or marginal asymmetries, but would not generally recover a distinction erased by the representation. The poor norm baseline is therefore partly an identifiability result, not merely a weak optimization outcome.

Temporal context might change that conclusion. A sequence of norms, its derivatives, or an asymmetric motion cycle could carry phase information absent from one instantaneous norm. The present baseline has no temporal window and should not be used to dismiss invariant sequence models. Likewise, a rotation-equivariant model could preserve directional relationships while accommodating coordinate changes. Those are plausible extensions with different information interfaces. Our experiment intentionally uses instantaneous samples so that the tradeoff between invariant compression and calibrated directional coordinates is easy to inspect.

Three known vectors also create a strong information advantage over raw axes. It would be misleading to present calibration as outperforming an equal-input competitor. The calibrated pipeline observes a separate reference experiment. Its benefit should be evaluated together with the cost and reliability of collecting that experiment. In a practical setting, a short calibration routine may be worthwhile, but that is an operational hypothesis requiring user compliance, repeatability, and sensor-error evaluation. None of those properties is measured by the generated reference vectors.

## 7 Where wearable evidence would begin

A real inertial signal includes gravity, translational acceleration, angular velocity, sensor bias, scale errors, and timing effects. Depending on the task, a coordinate transform alone may be insufficient. Sensor-to-segment motion can change during an activity, and soft-tissue movement can violate a fixed-frame assumption. The model study omits all these effects. It therefore offers a mechanism that can guide experimental design, while leaving clinical, sports, rehabilitation, and navigation performance entirely unvalidated.

The test phases are uniformly sampled over three repeated cycles with fixed amplitudes. Real motion can vary in speed, cycle shape, and amplitude, and event definitions may be uncertain. A future dataset should separate participants and sessions, preserve mounting metadata, and evaluate reference-pose errors independently of the activity predictor. It should also distinguish frame calibration from physiological label validity. A highly accurate synthetic half-cycle classifier cannot establish that a real foot-contact event has been measured correctly or that a downstream intervention is beneficial.

The graphics remain tied to the signal model. Coordinate triads state the transformation. Phase-colored loops show how shape is preserved while axes change. Circular links reveal angular disagreement without a misleading discontinuity at negative and positive pi. Trial distributions separate variability across rotations from the dense samples within one rotation. An athlete photograph or an anatomical figure would add application atmosphere but no evidence to this experiment. The original geometric visuals communicate the actual computation more precisely.

## 8 Coordinate information and phase information

### 8.1 A rotation changes coordinates without changing the loop

A proper rotation preserves lengths, angles, and the shape of the trajectory in three-dimensional space. It changes the numbers assigned to the axes. The nominal ridge predictor is tied to those axis numbers because its coefficients were fitted in the canonical frame. Applying it directly after rotation therefore composes the learned map with an unmodeled transformation. The resulting error can be large even when the underlying periodic signal is unchanged and the test phases remain inside the training range. This is a coordinate shift rather than a new activity distribution in the generator.

The distinction is useful for diagnosis. If a model fails after remounting, one should ask whether its input frame matches the frame assumed by training before concluding that the motion pattern has changed. In a real device, remounting can also alter sensor location, soft-tissue coupling, and dynamics, so a pure rotation may be insufficient. Our experiment isolates only the frame component. The large improvement after calibration shows what correcting that component can accomplish in a favorable setting, while leaving other mounting effects untested.

The trajectory figure preserves phase color across frames so that corresponding portions of the loop can be compared. The rotated loop is not a different generated activity. The calibrated loop approximately returns to the canonical coordinates, with deviations from reference noise and test noise. A reader can therefore see the input transformation that the predictor receives. This is more informative for the stated mechanism than a generic neural architecture diagram, which would not reveal whether its coordinate assumptions were satisfied.

### 8.2 Why the nominal map is adequate before rotation

The first canonical coordinate is 1.1 times sine phase and the second is 0.7 times cosine phase. Without noise, a linear map can recover the two target components exactly by rescaling those coordinates. The third coordinate adds a double-frequency term and an offset but is not necessary for that ideal recovery. The fitted ridge model is therefore not deliberately underpowered for the nominal task. Its failure after rotation is interpretable as a mismatch in coordinate meaning rather than an inability to represent phase in the original frame.

Noise and regularization prevent exact recovery in the executed fit. The coefficients are estimated from finitely many noisy samples, and the third coordinate can receive a small fitted contribution. The experiment does not claim a zero nominal phase error or compare its ridge fit with the analytical rescaling rule. A future baseline could include that exact map to isolate fitting error further. The current calibrated result already demonstrates that the learned nominal mapping remains useful once its input frame is approximately restored.

Training phases are independent uniform samples, while test phases are evenly spaced over repeated cycles. This changes the arrangement of samples but not the intended marginal coverage of phase. The predictor uses instantaneous vectors and has no temporal state, so it does not exploit the regular test ordering. The ordering is used for trajectory and within-cycle visualizations only. This distinction matters because a temporal model could infer phase progression from sequence structure even when an individual sample is ambiguous, which is outside the implemented task.

### 8.3 The norm has an exact half-cycle ambiguity

The canonical signal at phase plus pi reverses its first two coordinates and leaves its third coordinate unchanged. Its norm is therefore exactly the same. Yet both target components, sine and cosine, reverse sign. Under a uniform phase distribution, every norm-compatible phase has an equally likely opposite-phase partner with the opposite target. The conditional mean of the target components given norm is consequently zero in the ideal model. This is a stronger explanation for the poor norm baseline than merely saying that a fourth-degree polynomial is too simple.

$$\| a(\phi+\pi) \|_2 = \| a(\phi) \|_2 \tag{4}$$

The isotropic coordinate-noise law preserves this paired distributional symmetry. A half-turn about the third axis maps one noiseless vector to its opposite-phase counterpart and maps isotropic noise to a noise vector with the same distribution. Norm observations therefore remain unable to distinguish the paired half cycles under the stated generator. A more flexible instantaneous function of norm cannot systematically recover the missing sign information from independent samples. The learned polynomial's near-chance half-cycle result is consistent with that information limit.

The norm still varies with phase and can carry information about phase modulo a remaining ambiguity. It is not a useless summary for every task. A model predicting an unsigned periodic magnitude or a coarser invariant property could benefit from it. The present target specifically requires signed sine and cosine components, so the discarded direction matters. This illustrates a general representation principle: removing a nuisance transformation should be judged relative to the target, because the same invariant can be sufficient for one task and insufficient for another.

### 8.4 Temporal information would define a different experiment

A sequence of norms may contain derivatives, asymmetric timing, or continuity cues that help track phase. Our generator repeats a fixed loop at uniform phase speed, and the implemented norm model receives one sample at a time. It has no initial phase reference or temporal window. The impossibility statement therefore concerns the instantaneous norm under the specified symmetric prior, not every

sequence-based invariant representation. Adding history changes the information set and should be credited explicitly if it improves prediction.

A temporal model can also inherit an unresolved initial branch. If two phase trajectories differ by pi and produce the same norm sequence, continuity alone cannot identify which one is the canonical label without an initial cue or another asymmetric observation. The precise answer depends on the signal and target definitions. The current paper does not solve that sequence-identifiability problem, but the half-cycle symmetry suggests what an extension should test. A model's smooth output is not by itself evidence that the absolute phase branch has been observed.

Real movements often contain asymmetric phases and variable timing, which can break some of these symmetries. They can also introduce new nuisance variation that makes calibration harder. A practical study should evaluate the actual information in its signals rather than importing the exact norm limitation or the favorable calibration result from this generator. The analytic example is useful as a diagnostic boundary: it shows how a representation can erase a target distinction even when it successfully removes orientation.

### 8.5 What three reference vectors constrain

One labeled vector correspondence constrains a rotation only up to an arbitrary turn about that vector. Two noncollinear labeled directions can determine a proper rotation in the noise-free case, and a third provides additional geometric coverage. Our three canonical unit vectors are orthogonal, which is a favorable configuration. They constrain all axes rather than clustering near one direction. The small reference set is therefore informative because of its geometry and labels, not simply because it contains three numerical records.

If the reference vectors were nearly collinear, the estimated rotation could be sensitive around their shared direction. If their canonical values were uncertain, the Procrustes fit would inherit that uncertainty. If a label were swapped, the solution could be systematically wrong even with small measurement noise. None of those cases is included. The experiment assumes correct identities and known canonical vectors, then adds only independent isotropic measurement noise. Its calibration result is conditional on those strong but transparent reference assumptions.

A physical accelerometer observes gravity in static poses, which supplies a direction but not every orientation component from one pose alone. Other sensors or known rotations may be needed to establish a full frame. The generated unit-vector references should not be described as three arbitrary easy-to-collect wearable samples. They represent a controlled calibration experiment with known directions. A real protocol would need to show how those directions are established and how accurately a user or fixture can reproduce them.

### 8.6 The proper-rotation constraint matters

An unconstrained least-squares matrix could include scaling, shear, or reflection. Orthogonal Procrustes restricts the estimate to an orthogonal transform, and the determinant correction selects a proper rotation rather than a mirror transformation. This matches the test generator, which applies only proper rotations. The constraint uses valid structural knowledge and reduces the parameter family. It would be incomplete for a sensor with scale-factor error or cross-axis coupling, where a pure rotation cannot represent the full calibration map.

The singular-value decomposition solution is closed form for the stated correspondence objective, but this does not make every calibration problem closed form. Unknown correspondences, time offsets, biases, and nonlinear sensor errors can require different estimation procedures. We use the geometric solution because the experiment supplies exactly the information it requires. The resulting accuracy should not be credited to solving those unmodeled complications. A complete sensor-calibration claim would need to specify and validate the full error model.

The row-vector convention can be a source of implementation mistakes. Canonical rows are multiplied by the transpose of the generating rotation to form observed rows. Returning to canonical coordinates therefore requires right multiplication by the rotation itself. The estimated matrix is used in that direction. Applying its transpose instead could produce a plausible-looking but incorrect loop and a large phase error. The trajectory figure provides a visual check, while the mean rotation error provides a numerical check of the orientation estimate.

### 8.7 Drawing rotations without a preferred direction

A Gaussian matrix followed by QR factorization does not automatically yield the intended uniform rotation distribution unless sign conventions are handled. The triangular factor's diagonal signs are applied to the orthogonal factor, and a determinant adjustment restricts the result to proper rotations. This prevents a numerical convention from creating a preferred orientation pattern in the test set. The correction is part of the declared experiment rather than an optional cosmetic detail, because raw-axis performance can depend strongly on which rotations are sampled.

Forty rotations provide a finite sample of that distribution. Their average raw phase error need not equal a simple theoretical constant, particularly because the learned map and loop are not rotationally symmetric in their coordinate dependence. The observed 96.62 degrees should therefore be reported as the executed mean, not rounded into a claim that arbitrary rotation always produces exactly ninety-degree error. The norm result has a separate symmetry argument, while the raw result reflects how the fixed predictor interacts with the sampled frames.

A broader study could condition performance on rotation angle or axis, but such conditioning should be defined carefully. Rotation angle alone may not determine prediction difficulty because the loop and predictor privilege particular directions. A small rotation about one axis can affect phase differently from a larger rotation about another. The current empirical cumulative distribution shows trial variability without reducing it to an unsupported one-dimensional law. The first-trial trace illustrates one case and is not selected to represent every rotation.

### 8.8 Phase error is circular and task dependent

The two-argument arctangent maps predicted sine-cosine components to an angle, but its output has a numerical branch cut. Wrapped evaluation prevents a small difference across that cut from becoming an artificial large error. Component RMSE avoids the cut entirely but weights errors differently from angular distance. When the predicted vector has very small norm, its angle can be unstable even if component errors are moderate. This is especially relevant to the norm baseline, whose target-component predictions contain little signed information.

The half-cycle label is coarser. It asks only whether the sine component is positive. Errors near a true zero crossing can change that label with a small angular displacement, while a large displacement within the same half cycle can leave it correct. The high calibrated label accuracy and low phase error therefore support related but distinct conclusions. Neither metric is a measured gait event. The paper defines the label directly from generated phase so that its meaning remains mathematical rather than physiological.

Trial averaging also needs a unit. Each trial shares one rotation and one calibration estimate across six hundred samples. Those samples are not independent tests of calibration reliability. The paper first summarizes prediction within each trial and then averages across forty rotations. A confidence claim about new mounting orientations would need to use the trial hierarchy and perhaps many more reference realizations. Reporting twenty-four thousand scalar predictions as if they were independent calibration trials would overstate the evidence.

### 8.9 Calibration noise and prediction noise interact

The frame estimate is perturbed by reference noise, while each test vector has independent coordinate noise. Their effects on phase are not simply additive scalar angles. A frame error rotates the signal in a

direction that depends on phase, and the nominal predictor maps that perturbation through its coefficients. Test noise then adds sample-specific variation. This interaction explains why the mean phase error is not equal to the mean rotation error and why subtracting the two would not isolate a meaningful residual source.

The first-trial error trace shows local variation across phase even under a fixed rotation. That variation can arise because the sine-cosine target has different sensitivity to perturbations at different positions and because the third signal coordinate contains a second harmonic. The trace is diagnostic, not a fitted physical timing law. A formal sensitivity analysis could linearize the arctangent and predictor around each phase, but would need to handle small predicted norms and finite noise carefully. We do not use such a linearization to claim a global error bound.

Increasing reference repetitions could reduce rotation uncertainty under independent noise, but the benefit would eventually be limited by test noise and model fitting. Increasing training samples would address another component. A resource study could allocate effort among reference measurements, nominal training, and test sampling. The present experiment fixes all three. Its large calibration gain establishes the importance of the frame under this setup, while leaving the optimal calibration duration or measurement budget as an unexecuted design question.

### 8.10 The empirical next step has several boundaries

A wearable validation should record sensor placement, participant identity, session, reference-pose quality, and any remounting between calibration and activity. It should test new participants or sessions according to the intended use and should include realistic sensor errors. A model that succeeds when the same person repeats a controlled pose may not generalize to unconstrained field motion. The generated study contains no such hierarchy, so its result can motivate a protocol but cannot substitute for one.

Mounting drift is particularly different from a fixed trial rotation. If the sensor moves relative to the body during an activity, one initial calibration may become stale. A time-varying estimator or a model robust to relative motion could be needed. Likewise, a sensor moved to another body location experiences a different signal, not merely a rotated copy. These mechanisms should be separated experimentally. The current fixed-frame correction addresses one of them and should not be described as a complete solution to placement variability.

The most defensible transfer is therefore methodological. State the coordinate frame assumed by the predictor, identify what calibration information is available, and test whether an invariant representation retains the target distinction. Use circular metrics for phase and respect trial-level dependence. The original triads, loops, phase links, and error distributions make those requirements visible in one transparent model. They provide a concrete baseline for later sensor studies while preserving the boundary between a mathematical signal experiment and evidence about people or devices.

### 8.11 The comparison changes access as well as representation

The calibrated pipeline receives three labeled reference vectors unavailable to the raw and norm pipelines. This is the intended intervention, and it should be visible whenever the result is summarized. Calling the methods equal-input competitors would conceal the most important reason for improvement. A fair practical comparison could charge the reference acquisition cost or give other models the same references in another form. The current study instead measures the value of that added frame information under a simple geometric solution.

The norm pipeline uses a different ridge penalty and feature family from the axis pipeline because its powers have different scales and a different intended invariance. It is a diagnostic representation baseline, not a carefully optimized competitive model. Its exact half-cycle ambiguity explains why greater instantaneous flexibility would not solve the signed target under this generator. Nevertheless, a

claim about the best practical invariant model would require temporal features, validation, and realistic signals. We do not make that claim from one fourth-degree polynomial.

The raw and calibrated axis predictions do use the same nominal coefficients. That pairing isolates the coordinate correction more cleanly: the predictor is not retrained with test phase labels, and its capacity is unchanged. The calibration estimate changes only the input frame. This is a useful controlled design because it separates a geometrical intervention from a learning improvement. It also shows why a simple predictor can appear inadequate when an upstream representation assumption is violated, then become useful again after that assumption is restored.

### 8.12 A reference routine needs its own validation

A physical reference routine can fail through incorrect pose execution, mislabeled directions, sensor movement, or a mismatch between the calibration and activity frame. These errors may be systematic rather than independent Gaussian noise. Repeating a flawed routine can reduce random variation while preserving bias. The generated three-vector set includes none of those complications, so the small rotation error should not be interpreted as evidence that users can achieve the same calibration quality. The routine itself would need an empirical repeatability study.

Reference quality could be checked by holding out an additional known direction and measuring its alignment after calibration. That would test the fitted frame on information not used in the Procrustes solve. It would still not validate the activity model or compensate for a sensor moved to a different physical location. The current experiment uses only three references and does not report such a held-out-direction test. The proposal identifies a concrete next check without implying that it has already been passed.

A calibration that remains stable over time would also need to be distinguished from one that is valid only immediately after mounting. Temperature, mechanical drift, and repeated motion can alter sensor behavior. The test trials hold rotation constant and use independent stationary noise, so they cannot assess persistence. A future study should repeat reference checks before and after activity and should evaluate whether the correction remains appropriate. These measurements would extend the evidence from a coordinate mechanism to a usable sensing protocol.

### 8.13 A phase model should state its anchor

Absolute phase requires a convention about where the cycle begins. In this generator, that convention is built into sine and cosine. In a real activity, a phase anchor might be a physical event, an external measurement, or a learned annotation. Its uncertainty can affect every phase label even if coordinate calibration is perfect. The current study fixes the anchor and tests orientation only. A human-motion claim would need to validate the anchor separately rather than assuming that a mathematically smooth phase corresponds to a meaningful event.

The same distinction applies to frequency and cycle duration. Test samples cover three uniform cycles, so the predictor is not asked to estimate a changing cadence or detect cycle boundaries. A temporal system might need to solve those tasks jointly. The present instantaneous mapping can be evaluated independently because generated phase is known to the evaluator. That is a useful experimental simplification, but it prevents the reported accuracy from being translated directly into timing performance during irregular sports or rehabilitation movement.

The final interpretation is therefore deliberately layered. The nominal signal contains enough directional information for a simple phase map. Arbitrary rotation changes that coordinate meaning. Norm compression removes the rotation and a necessary sign distinction together. Labeled references recover the frame and preserve the target information. Each step is supported by the generator, algebra, and executed predictions. Real sensor and human evidence would add new layers, whose success should be measured rather than inferred from the clarity of this mathematical example.

## 9 Conclusion

A fixed but unknown sensor rotation can invalidate an otherwise adequate instantaneous phase predictor. In the generated study, three labeled reference vectors restore the coordinate frame and greatly reduce phase error, while a norm-based invariant discards necessary directional information. The result supports evaluating nuisance invariance together with target sufficiency and calibration access. It does not establish a wearable device, a gait model, or human performance. Reference uncertainty, temporal context, mounting drift, and real sensor errors are the next empirical questions.

# References


[1] R. Mahony, T. Hamel, J. M. Pflimlin, “Nonlinear Complementary Filters on the Special Orthogonal Group,” IEEE Transactions on Automatic Control, vol. 53, pp. 1203-1218, 2008. doi: 10.1109/tac.2008.923738.

[2] S. O. H. Madgwick, A. J. L. Harrison, R. Vaidyanathan, “Estimation of IMU and MARG orientation using a gradient descent algorithm,” 2011 IEEE International Conference on Rehabilitation Robotics, pp. 1-7, 2011. doi: 10.1109/icorr.2011.5975346.

[3] A. M. Sabatini, “Quaternion-Based Extended Kalman Filter for Determining Orientation by Inertial and Magnetic Sensing,” IEEE Transactions on Biomedical Engineering, vol. 53, pp. 1346-1356, 2006. doi: 10.1109/tbme.2006.875664.

[4] T. Seel, J. Raisch, T. Schauer, “IMU-Based Joint Angle Measurement for Gait Analysis,” Sensors, vol. 14, pp. 6891-6909, 2014. doi: 10.3390/s140406891.

[5] I. H. Lopez-Nava, A. Munoz-Melendez, “Wearable Inertial Sensors for Human Motion Analysis: A Review,” IEEE Sensors Journal, vol. 16, pp. 7821-7834, 2016. doi: 10.1109/jsen.2016.2609392.

[6] L. Pacher, C. Chatellier, R. Vauzelle et al., “Sensor-to-Segment Calibration Methodologies for Lower-Body Kinematic Analysis with Inertial Sensors: A Systematic Review,” Sensors, vol. 20, pp. 3322, 2020. doi: 10.3390/s20113322.

[7] F. Ordóñez, D. Roggen, “Deep Convolutional and LSTM Recurrent Neural Networks for Multimodal Wearable Activity Recognition,” Sensors, vol. 16, pp. 115, 2016. doi: 10.3390/s16010115.

[8] J. Wang, Y. Chen, S. Hao et al., “Deep learning for sensor-based activity recognition: A survey,” Pattern Recognition Letters, vol. 119, pp. 3-11, 2019. doi: 10.1016/j.patrec.2018.02.010.

[9] L. Bao, S. S. Intille, “Activity Recognition from User-Annotated Acceleration Data,” Lecture Notes in Computer Science, pp. 1-17, 2004. doi: 10.1007/978-3-540-24646-6_1.

[10] W. Kabsch, “A solution for the best rotation to relate two sets of vectors,” Acta Crystallographica Section A, vol. 32, pp. 922-923, 1976. doi: 10.1107/s0567739476001873.

[11] S. Umeyama, “Least-squares estimation of transformation parameters between two point patterns,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 13, pp. 376-380, 1991. doi: 10.1109/34.88573.

[12] S. Patel, H. Park, P. Bonato et al., “A review of wearable sensors and systems with application in rehabilitation,” Journal of NeuroEngineering and Rehabilitation, vol. 9, pp. 21, 2012. doi: 10.1186/1743-0003-9-21.

[13] X. Wang, H. Yu, S. Kold et al., “Wearable sensors for activity monitoring and motion control: A review,” Biomimetic Intelligence and Robotics, vol. 3, pp. 100089, 2023. doi: 10.1016/j.birob.2023.100089.

[14] Y. Suh, “Inertial Sensor-Based Smoother for Gait Analysis,” Sensors, vol. 14, pp. 24338-24357, 2014. doi: 10.3390/s141224338.

[15] D. Ravi, C. Wong, B. Lo et al., “Deep learning for human activity recognition: A resource efficient implementation on low-power devices,” 2016 IEEE 13th International Conference on Wearable and Implantable Body Sensor Networks (BSN), pp. 71-76, 2016. doi: 10.1109/bsn.2016.7516235.

[16] L. Breiman, “Random Forests,” Machine Learning, vol. 45, pp. 5-32, 2001. doi: 10.1023/a:1010933404324.

[17] C. M. Bishop, “Pattern Recognition and Machine Learning,” Springer, 2006. doi: 10.1007/978-0-387-45528-0.

[18] Y. Yu, T. Huang, Y. Yang, “Attitude-Aided Linear Calibration of Triaxial Accelerometers,” arXiv preprint arXiv:2606.06308, 2026. [Online]. Available: https://arxiv.org/abs/2606.06308.