Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A measurement protocol and falsifiable predictions for characterizing acoustic injection vulnerability of MEMS inertial sensors

FU SEN

fs211314@gmail.com

Changsha Vocational and Technical College

TAN SHUXIN

Hunan Women'S University

Article

Keywords: MEMS inertial sensors, acoustic injection, measurement protocol, falsifiable prediction, sensor characterization, UAV security

Posted Date: November 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11525835/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

A measurement protocol and falsifiable predictions for characterizing acoustic injection vulnerability of MEMS inertial sensors

Sen Fu[1*] and Shuxin Tan[2]

[1]Changsha Vocational & Technical College, Changsha, China.
[2]Hunan Women's University, Changsha, China.

*Corresponding author(s). E-mail(s): fs211314@gmail.com;
Contributing authors: tanyatsx92@163.com;

Abstract

Acoustic excitation can corrupt microelectromechanical-system (MEMS) gyroscopes and accelerometers, yet published results report incompatible quantities and therefore cannot be compared or pooled. We argue that the bottleneck is comparability rather than existence, and we contribute a measurement protocol rather than a new measurement. We derive five quantities that constitute a minimum characterization basis for comparing vulnerability, each with a measurement procedure, reporting units and required precision, and we state four falsifiable predictions with explicit refutation conditions, including a test that separates mechanical resonance from electronics-mediated down-conversion. We define a minimum reportable unit of nine items. Applying the protocol to two published studies shows that the ranking between them reverses in 34.6% of draws while each index spans about two decades, so any single assumed-parameter ranking is unstable; supplying the protocol's quantities removes this ambiguity. Computational tests confirm the predictions when the model holds and reject them in four deliberately constructed failure modes. No physical acoustic measurement is reported: the contribution is a specification that converts an uncomparable literature into comparable measurements.

Keywords: MEMS inertial sensors, acoustic injection, measurement protocol, falsifiable prediction, sensor characterization, UAV security

1 Introduction

Small unmanned aerial vehicles (UAVs) rely on low-cost MEMS inertial sensors to close attitude and navigation loops. This dependence makes the inertial measurement unit (IMU) a physical attack surface: an acoustic wave can couple into a vibrating proof mass, appear as a false inertial signal, and enter the flight controller before higher-level navigation sensors reject it. The effect is well established. Intentional sound disturbs gyroscopic sensors on drones [1]; acoustic injection can bias or control capacitive accelerometer outputs [2]; targeted ultrasonic excitation degrades gyroscopes [3]; and signal recovery methods have been built for compromised drone IMUs [4]. Related MEMS work establishes why proof-mass structures and capacitive readout respond to external vibration and why geometry governs sensitivity [5–9], and platform-level studies show that mobility and motor-induced frame vibration change outcomes relative to idealized settings [10, 11].

The problem this paper addresses is different from showing that the effect exists. It is that published results are not comparable. One study reports sound pressure level at a stated distance, another reports induced drift in deg/s, a third reports whether a drone crashed, and a fourth reports a resonant frequency for a specific axis of a named part. None of these can be converted into another without quantities that are almost never reported: the acoustic-to-output gain in physical units, the damping ratio of the excited mode, the attenuation introduced by the airframe and the incidence angle, and the rejection of the analog front end at the controller sampling rate. The practical consequence is that the same sensor physics can be described as severe in one paper and marginal in another, with no way to tell whether the disagreement is real or an artifact of what was measured.

This matters most for counter-UAV assessment, where the question is feasibility rather than existence. Deciding whether acoustic resonance is a dependable interdiction mechanism requires an envelope over frequency, pressure, damping and coupling, not a single demonstration. A protocol is what makes such an envelope accumulable across laboratories.

We state the positioning explicitly, because it determines how the work should be judged. This is a sensor characterization methodology, not a hardware-development or experimental-sensor study. **No physical acoustic measurement is reported.** The computational results characterise the behaviour of the protocol, not the behaviour of any device.

Acoustic environmental testing is not unstandardized. Military and industrial standards specify acoustic exposure tests for equipment, and metrological standards govern microphone calibration and sensor performance definition. These are *compliance* instruments: they define an exposure profile and a pass or fail criterion, and ask whether the unit continues to meet its specification. They are not designed to produce a result convertible into another laboratory's result, because the quantities they require are chosen for survivability assessment rather than cross-study comparability. The protocol here is therefore not a replacement for compliance testing; it is an additional reporting layer to be applied when the purpose is to characterize and compare susceptibility, and its procedures can be executed inside a standard acoustic test campaign.

Our contribution is fourfold: (i) five quantities that provide a minimum characterization basis for comparing vulnerability under the stated modelling assumptions, each with a procedure, units and required precision; (ii) four falsifiable predictions with explicit refutation conditions, assigned to the level that generates them; (iii) a minimum reportable unit of nine items; and (iv) a demonstration that the protocol discriminates and can fail.

2 Results

2.1 A reduced-order model and what it makes observable

We derive the model only to the extent needed to justify the protocol; it is not offered as a validated physical description. Let $q(t)$ be the generalized displacement of one sensitive axis and let the incident pressure be $p(t) = P\cos(\omega_a t + \phi)$. The dominant modal coordinate obeys

$$m\ddot{q} + c\dot{q} + kq = \alpha_p A_e p(t) + d(t), \tag{1}$$

with modal mass, damping and stiffness m, c, k, effective acoustic area A_e, packaging coupling α_p, and non-acoustic vibration $d(t)$. With $\omega_n = \sqrt{k/m}$ and $\zeta = c/(2\sqrt{km})$,

$$|H_m(j\omega_a)| = \frac{\alpha_p A_e/k}{\sqrt{(1-r^2)^2 + (2\zeta r)^2}}, \quad r = \frac{\omega_a}{\omega_n}. \tag{2}$$

For a vibratory gyroscope the acoustic displacement perturbs the driven mode and appears as apparent rate after demodulation,

$$\tilde{\Omega}_i(t) = K_{g,i}\,\mathcal{D}_i\{q_i(t)\} + \eta_{g,i}(t), \tag{3}$$

with readout gain $K_{g,i}$, demodulation and leakage operator $\mathcal{D}_i$, and nominal noise $\eta_{g,i}$. For a capacitive accelerometer, $\tilde{a}_i(t) = K_{a,i}q_i(t) + \eta_{a,i}(t)$.

Writing the false output as the two-path form of Eq. (7) shows why conventional reporting underdetermines vulnerability: a report of sound pressure level gives P but not η_f or κ_i; a report of induced drift gives the product but not its factors; a report of a crash gives neither.

2.2 Frequency conversion and sampling aliasing are distinct

A linear system driven at a single frequency responds at that same frequency; it does not generate a difference frequency. We verified this directly: evaluating the response of Eq. (1) to a single tone yields spectral content at the drive frequency and at the transient natural response, with no component at $|f_a - f_n|$. A difference frequency therefore requires a second sustained oscillation and a multiplication. In a vibratory gyroscope the second oscillation is the drive carrier f_d, and the multiplication is performed by demodulation, already represented by $\mathcal{D}_i$ in Eq. (3). The conversion product is

$$\Delta f_{\mathrm{mix},i} = |f_a - f_{d,i}|. \tag{4}$$

We further verified that a single-tone quadratic mechanical term βq^2 does not produce $|f_a - f_n|$, since q^2 generates a static term and a second harmonic rather than a difference tone; a difference frequency requires two sustained tones. Eq. (4) is therefore grounded in demodulation against the drive carrier.

It is important to be precise about which subsystem Eq. (4) describes. It is not a prediction of the mechanical subsystem governed by Eq. (1), which is linear and frequency preserving. It is a prediction about the complete sensor readout chain, comprising the mechanical response, the demodulation operator $\mathcal{D}_i$, and the baseband stage that follows. Nothing in the mechanical equation alone implies a converted frequency; conversion arises where the mechanical response is multiplied by the drive carrier during readout. Eq. (4) should be read as a readout-chain relation under mode-matched demodulation, and not as a mechanical nonlinearity.

Assumption 1 (Mode-matched operation)**.** *The drive or demodulation carrier satisfies $f_{d,i} \approx f_{n,i}$, as in mode-matched or near-mode-matched designs. Under this assumption $\Delta f_{\mathrm{mix}} \approx |f_a - f_n|$. When the design is mode-split, f_d must be measured and reported separately.*

The baseband stage passes the conversion product according to

$$|H_{e,i}(\Delta f)| = \frac{1}{\sqrt{1 + (\Delta f / f_{c,i})^4}}. \tag{5}$$

A second and distinct mechanism is sampling aliasing. Imperfect stop-band rejection leaks a fraction ε of the carrier, and analog-to-digital conversion folds it to

$$\Delta f_{\mathrm{alias},i} = \big| f_a - \mathrm{round}(f_a / f_s) f_s \big|, \tag{6}$$

which is an analog-to-digital sampling effect and not a physical mixing process. We emphasise the distinction because Δf_{mix} is independent of f_s whereas $\Delta f_{\mathrm{alias}}$ is independent of f_d, which is what allows a later study to attribute an observed low-rate disturbance to one or the other.

Combining both paths, and bounding the result by amplifier saturation,

$$\begin{aligned} \tilde{y}_i[k] &= \eta_f P \, \kappa_i \big(H_{m,i} H_{e,i}^{\mathrm{mix}} + \varepsilon H_{e,i}^{\mathrm{al}} \big), \\ y_i[k] &= \mathrm{sat}_{Y_i}\big(\tilde{y}_i[k]\big) + \nu_i[k], \end{aligned} \tag{7}$$

where $H_{e,i}^{\mathrm{mix}} = |H_{e,i}(\Delta f_{\mathrm{mix}})|$, $H_{e,i}^{\mathrm{al}} = |H_{e,i}(\Delta f_{\mathrm{alias}})|$, κ_i carries units of $(^\circ/\mathrm{s})/\mathrm{Pa}$ for gyroscopes and $(\mathrm{m/s^2})/\mathrm{Pa}$ for accelerometers, and $\mathrm{sat}_Y(x) = Y \tanh(x/Y)$.

Because flight stacks use both sensor types, define peak false magnitudes over a tested band $\mathcal{F}$ as

$$G(P) = \max_{i \in \mathcal{G}} \max_{f_a \in \mathcal{F}} |y_{g,i}|, \qquad A(P) = \max_{i \in \mathcal{A}} \max_{f_a \in \mathcal{F}} |y_{a,i}|, \tag{8}$$

and, with application tolerances T_g and T_a,

$$\mathcal{V}(P) = \left[\left(\frac{G(P)}{T_g} \right)^2 + \left(\frac{A(P)}{T_a} \right)^2 \right]^{1/2}. \tag{9}$$

Values above unity mean the disturbance requires rejection, recovery or explicit controller tolerance; they do not by themselves imply loss of control. The aggregation in Eq. (9) is a reporting choice rather than a uniquely implied physical law. We use the Euclidean form because it is monotone in each channel, reduces correctly to a single-channel index when the other is negligible, and corresponds to independent contributions to a quadratic risk of the kind an estimator approximately minimizes. A maximum or an arithmetic mean would also be defensible and would give different values, so the aggregation and the tolerances must be reported together with the index.

Proposition 1. *For fixed electronics and pressure, the mechanical factor in Eq. (2) is locally maximized near* $\omega_a = \omega_n\sqrt{1-2\zeta^2}$ *for* $\zeta < 1/\sqrt{2}$*, with peak magnitude scaling approximately as* $1/(2\zeta)$*.*

This is a standard property of a lightly damped second-order system rather than a new result; it is included because it generates two of the protocol's predictions.

2.3 The protocol: five required quantities

2.3.1 (M1) Modal frequency and damping

Definition. Natural frequency f_n and damping ratio ζ (equivalently $Q = 1/(2\zeta)$) of each sensitive axis. *Procedure.* Low-level frequency sweep with the device stationary, driven far below saturation, using acoustic excitation at fixed low pressure or mechanical shaker excitation; extract f_n from the response peak and ζ from the half-power bandwidth, $\zeta \approx \Delta f_{-3\mathrm{dB}}/(2f_n)$. *Units and precision.* Hz and dimensionless; $\pm 1\%$ on f_n and $\pm 10\%$ on ζ, which is what is needed to test **(P1)** and **(P2)**. *Common error.* Sweeping only until a response is seen and then reporting one resonance for the whole device; axes resonate at different frequencies, so per-axis values are required.

2.3.2 (M2) Effective acoustic-to-output transfer gain

Definition. The transfer gain measured directly at the terminals,

$$K_i(f) = \frac{Y_i(f)}{P(f)}, \tag{10}$$

in (°/s)/Pa for gyroscopes and $(\mathrm{m/s^2})$/Pa for accelerometers. *Procedure.* Place the device in a calibrated free field with a reference microphone at the device location; sweep pressure at fixed frequency inside the linear range and regress output amplitude against pressure in Pa; report the curve $K_i(f)$ rather than a single number. *Relation*

to the latent factor. What is measured is the product, not the isolated factor:

$$K_i(f) = \kappa_i \, |H_{m,i}(f)| \, |H_{e,i}(\Delta f)|. \tag{11}$$

The latent factor κ_i is obtainable only when $|H_m|$ and $|H_e|$ have been characterized independently through **(M1)** and **(M4)**. *Common error.* Reporting sound pressure level in dB without a reference distance and microphone calibration, or reporting amplitude in arbitrary units or dimensionless codes; this is the most frequent omission in the literature we reviewed.

2.3.3 (M3) Incidence and frame attenuation

Definition. $\eta_f \in (0, 1]$, the factor by which airframe, packaging and incidence reduce the effective pressure at the die. *Procedure.* Rotate the device or airframe on a calibrated turntable in a free field and record response versus incidence angle; report bare-sensor and mounted values separately so packaging and airframe contributions can be separated. *Common error.* Measuring the bare sensor and generalizing to the aircraft; this is the step that most often turns a laboratory result into an incorrect operational claim.

2.3.4 (M4) Baseband rejection and sampling residue

Definition. The front-end magnitude response $|H_e|$ at baseband, the controller sampling rate f_s, the anti-alias cutoff, and the leakage floor ε. *Procedure.* Inject known tones across the band and measure the residual amplitude appearing at the controller output rate; report f_s, filter order and cutoff, and measured clock jitter. *Precision.* Because $\Delta(\Delta f_{\text{alias}}) = k \, \Delta f_s$ with $k = \text{round}(f_a / f_s)$, the alias position error is *amplified* by k; achieving a 1 Hz alias-position accuracy at $k = 7$ requires $\Delta f_s \lesssim 0.14$ Hz, so a single fixed percentage is inadequate and the requirement must be stated per operating point. *Common error.* Assuming that because a tone is above Nyquist it cannot affect the output; Eq. (6) shows why that inference fails.

2.3.5 (M5) Application-specific disturbance tolerance

Definition. Two quantities T_g and T_a, the false gyroscope and accelerometer magnitudes at which the consuming estimator or controller ceases to reject the disturbance. *Procedure.* Where the flight stack is known, T_g and T_a are its rejection limits; where no specific stack is available, record instead the nominal disturbance reference

$$Q_{0.99} = 99\text{th percentile of nominal innovation magnitude}, \tag{12}$$

and label it a *baseline reference*, not a tolerance. We insist on this distinction because a percentile of ordinary operation describes the noise floor whereas a tolerance is a property of the controller; conflating them produces an index that looks interpretable but is not.

Table 1 summarises the five quantities.

Table 1 The five required quantities. For **(M2)** the directly measured quantity is the transfer gain $K_i(f)$; the factor κ_i is a de-embedded latent quantity available only once **(M1)** and **(M4)** are independently characterized.

ID	Quantity	Units	Precision
M1	f_n, ζ per axis	Hz, –	$\pm 1\%$, $\pm 10\%$
M2	$K_i(f)$ (direct); κ_i latent	out/Pa	slope + range
M3	η_f per incidence	–	per axis
M4	H_e, f_s, ε	Hz, –	per operating point
M5	T_g, T_a (or $Q_{0.99}$ reference)	°/s, m/s^2	quantile

2.4 Falsifiable predictions

A protocol is only useful if it can be wrong. We state four predictions, each with the measurement needed and the result that would refute it.

(P1) *Peak position is locked to the mode.* The false-output maximum lies at $f_n\sqrt{1-2\zeta^2}$ and does not move with acoustic pressure. Requires **(M1)**, **(M2)**. *Refuted if* the peak position shifts systematically with pressure beyond the precision of **(M1)**, which would indicate amplitude-dependent stiffening or a nonlinearity outside the model.

(P2) *Peak amplitude scales inversely with damping.* At fixed pressure and electronics the peak scales approximately as $1/(2\zeta)$. Requires **(M1)**, **(M2)** and at least two devices or modes differing in ζ. *Controlled condition.* **(P2)** is a ceteris-paribus prediction: damping does not vary independently in a real device, and a change in ζ is typically accompanied by changes in gain, bandwidth, mode coupling and front-end response. The prediction attributes an amplitude change to damping only if the remaining transfer-chain parameters are held constant or characterized independently through **(M1)**, **(M2)** and **(M4)**. Where they are not, an observed amplitude change is consistent with the damping law but does not establish it.

(P3) *Conversion versus aliasing.* When the injected tone lies above Nyquist, the disturbance appearing at the controller rate is at either the conversion product $|f_a - f_d|$ or the alias image $|f_a - \mathrm{round}(f_a/f_s)f_s|$. Requires **(M1)**, **(M4)**. *Refuted if* the observed baseband component does not coincide, within the cutoff bandwidth, with either predicted position. *Not refuted by waveform shape alone*: a linear system driven by a single sustained tone produces a sinusoid at the converted frequency, so a sinusoidal appearance carries no information about whether the component arises from conversion, from aliasing, or from a synthetically injected signal. Distinguishing synthetic injection from physical response instead requires a dependence test, namely that the component scale with measured pressure and vary with measured incidence and mounting, as specified by **(M2)** and **(M3)**.

(P4) *Attenuation changes the index monotonically.* $\mathcal{V}$ decreases monotonically as η_f decreases across incidence and mounting configurations. Requires **(M2)**, **(M3)**, **(M5)**. *Refuted if* $\mathcal{V}$ is insensitive to mounting and incidence; insensitivity would indicate that the assumed acoustic attenuation path is insufficient to explain the observation and would motivate examination of alternative coupling mechanisms, including

Table 2 Level, required measurement, and refutation condition for each prediction.

Pred.	Level	Requires	Refuted if
P1	mechanical	M1, M2	peak drifts with P
P2	mechanical	M1, M2	amplitude insensitive to ζ
P3	readout chain	M1, M4	baseband off-image
P4	reporting convention	M2, M3, M5	index insensitive to η_f

Table 3 Competing accounts and their opposing predictions.

Observation	Resonance	Electronics-only
Peak location	tracks f_n	tracks f_s comb
Bandwidth	narrow, set by ζ	broad
Sensitivity to damping	strong, $1/(2\zeta)$	weak
Response to detuning	peak moves	peak fixed

structural conduction and electromagnetic coupling. We deliberately do not read insensitivity as positive evidence for any particular alternative, since a single monotonicity test cannot identify which additional path is present.

Table 2 assigns each prediction to the level that generates it, because conflating levels invites the objection that a single negative result refutes the whole protocol. **(P1)** and **(P2)** are mechanical predictions following from Eq. (2); **(P3)** is a readout-chain prediction following from Eqs. (4) and (6) and is not implied by the mechanical equation; **(P4)** is, strictly, a consequence of the reporting convention in Eq. (9) together with the linearity of η_f in the unsaturated regime, rather than a claim about sensor physics. The distinction matters for attribution: a failure of **(P3)** localises the problem to the readout or sampling description, whereas a failure of **(P1)** localises it to the mechanical description.

Table 3 states what each account predicts; Fig. 1 shows the resulting difference on the synthetic axis.

2.5 Minimum reportable unit

We define the minimum reportable unit as nine items to be reported together for every vulnerability claim: sensor make and model; axis; f_n and ζ per axis (**(M1)**); K_i with its linear range (**(M2)**); η_f with mounting description (**(M3)**); f_s, anti-alias cutoff and ε (**(M4)**); T_g and T_a with quantile (**(M5)**); the frequency and pressure grid actually sampled; and the resulting $\mathcal{V}$ or equivalent index. Two reporting rules follow. First, report products and factors: induced drift alone is not sufficient, and neither is sound pressure level alone. Second, report negative and null results with the same

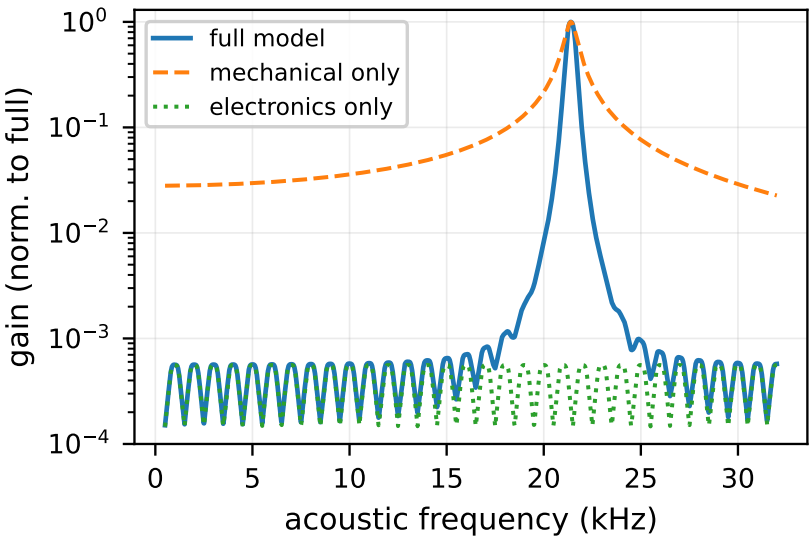


Fig. 1 Discriminative comparison on the synthetic 21.4 kHz axis. The full model is narrow and locked to f_n; the electronics-only account is broad and locked to the sampling grid.

fields, since a null obtained at a stated pressure with a stated η_f bounds feasibility and is more useful than a positive result with no context.

2.6 Worked comparability example

The protocol's central claim is that supplying **(M1)** through **(M5)** makes results comparable, so that claim needs an explicit test. We use two published sources and distinguish carefully what each reports from what we must assume.

Study A is the WALNUT investigation of capacitive MEMS accelerometers [2], which reports a resonance at 2.9 kHz for the ADXL337 and the observation of output biasing. Study B is the RAiS dataset [12], which reports a Z-axis gyroscope resonance at 27.55 kHz for a VectorNav VN100, a sampling rate of 40 Hz, and severity bands in which the hard band exceeds 6 °/s of induced drift. The two report no quantity in common apart from a resonant frequency, and they concern different sensor types.

We normalize both to a common incident pressure $P = 20$ Pa. At resonance, with the converted component inside the baseband cutoff, Eq. (9) reduces to

$$\mathcal{V} \approx \frac{\eta_f \, P \, \kappa}{2\zeta \, T}, \tag{13}$$

where T is the applicable tolerance. Every quantity in Eq. (13) except f_n is unreported by at least one study, so we draw the unreported quantities independently from ranges typical of MEMS inertial sensors: $\zeta \in [0.005, 0.05]$, $\kappa \in [0.001, 0.10]$ $(\mathrm{m/s^2})/\mathrm{Pa}$ for the accelerometer and $[0.01, 1.00]$ $(^\circ/\mathrm{s})/\mathrm{Pa}$ for the gyroscope, $\eta_f \in [0.2, 1.0]$, $T_a \in [0.5, 5.0]$ $\mathrm{m/s^2}$ and $T_g \in [10, 100]$ $^\circ/\mathrm{s}$, over 20,000 draws.

Study A yields the larger index in 65.4% of draws and study B in the remaining 34.6%. Each study's own index spans 1.87 and 1.86 decades respectively, so the absolute level is undetermined by nearly two orders of magnitude purely as a consequence of unreported quantities. Within any single draw the comparison often looks decisive: the ratio exceeds two or falls below one half in 72.4% of draws.

Several assumed ranges span one to two decades, so the distribution chosen over them is not innocuous and we tested it. Repeating the analysis with log-uniform draws, the more natural default when a quantity is known only to within an order of magnitude, gives 61.5% versus 38.5% and a decisive fraction of 79.0%, with spreads of 2.41 decades for both studies. The reversal fraction moves by only about four percentage points between the two distributional assumptions, so the conclusion is stable. We state explicitly that these fractions are not population probabilities and must not be read as the chance that one device is truly more susceptible than the other, since the distributions are envelopes specified by us rather than measured dispersions; they are sensitivity measures quantifying how much of the comparison is determined by quantities the source studies did not report.

The combination of these facts is the finding. An analyst who supplies a single plausible set of assumed parameters obtains a confident-looking ranking, but the direction of that ranking is determined by quantities the source studies never reported, and a different plausible assumption reverses it about one time in three. Incomparability of this kind is dangerous precisely because it does not announce itself. We stress that the example does *not* show that the ADXL337 is more susceptible than the VN100 gyroscope; it shows that the published content of the two studies cannot support any ranking.

Which quantities each study reports and which it omits is recorded in Supplementary Table S7. Repeating the comparison with **(M1)** through **(M5)** supplied as point values removes the ambiguity entirely: the index evaluates to 1.79 for study A and 0.89 for study B, and the ordering is no longer draw-dependent.

2.7 Computational verification

We demonstrate that the protocol is executable and that its tests can both succeed and fail. We deliberately avoid the word validation: everything here is computational, the parameters are assigned rather than measured, and the results characterise the protocol, not any sensor.

We instantiate one gyroscope-like channel with $f_n = 21.4\,\mathrm{kHz}$, $\zeta = 0.014$, $f_d = f_n$ under Assumption 1, $f_c = 260\,\mathrm{Hz}$, $f_s = 1000\,\mathrm{Hz}$, $\varepsilon = 0.02$, $\kappa = 0.13$ (°/s)/Pa and $P = 20\,\mathrm{Pa}$.

(P1) holds: varying f_n over 16–27 kHz, the located peak tracks the assigned mode to within 0.004 kHz, far inside the 1% required by **(M1)**. **(P2)** holds: over an eightfold damping range the product of peak amplitude and ζ is constant to a spread of 1.051, the residual arising from the leakage path (Fig. 2b). **(P3)** holds: varying f_s from 500 to 4000 Hz leaves the conversion product fixed while the alias image moves, and the two predicted positions differ by more than the cutoff bandwidth in every configuration. **(P4)** holds: in the unsaturated regime the index is exactly proportional to η_f, falling

monotonically from 1.223 to 0.383 as η_f goes from 1.0 to 0.31, a ratio of 0.313 against an expected 0.310. Full values are in Supplementary Tables S1–S4.

Noise and parameter uncertainty. We add white noise and measure whether the mode can still be localized, defining SNR against the resonant peak amplitude, which is the quantity actually detected. An earlier version referenced the noise to the signal standard deviation and consequently reported successful detection at every level including 5 dB; that was an artifact of the reference choice and is corrected here. With the peak-referenced definition the detection rate is 1.000 at 40, 30 and 20 dB, 0.995 at 10 dB and 0.660 at 5 dB (Fig. 2c), giving a usable operating floor near 10 dB. Under $\pm 10\%$ independent dispersion of f_n, ζ, κ and η_f over 3000 draws the median normalized disturbance is 1.225 with a 5–95% interval of 1.035–1.434, so configuration ranking is stable while the absolute value carries roughly $\pm 17\%$ spread (Supplementary Table S5).

Failure-mode tests. A falsifiable protocol must be able to fail. We construct four cases in which the underlying assumption is deliberately violated, and each corresponding prediction is refuted: a mis-assigned mode leaves a 2.6 kHz residual (**(P1)**); an unchanged amplitude under doubled ζ gives a ratio of 1.00 against an expected 0.50 (**(P2)**); a baseband component 150 Hz off the predicted alias falls outside the cutoff band (**(P3)**); and an unchanged index under halved η_f gives 1.00 against 0.50 (**(P4)**). Details are in Supplementary Table S6. These cases are the main evidence that the predictions carry information, since a test that cannot fail does not constrain anything.

External resource audit. We treat external data as a provenance audit rather than corroboration, and state why. Applying **(M4)** and **(P3)** to the IRIS_HITL dataset released with UNROCKER [4, 13], which states an injection at 1830 Hz and a controller rate of 250 Hz, Eq. (6) predicts an image at $|1830 - 7 \times 250| = 80$ Hz and the dominant component is at 80.23 Hz, a difference of 0.23 Hz. The released resource performs software-only injection: we verified that the difference between the contaminated channel `attacc_accel0` and the reference `true_accel0` is identical to the resource-written `attacc_offset` column, with a Pearson correlation of exactly 1.000000 over all 33,153 samples. The component is therefore the output of an independently written model and not a physically measured alias; the 0.23 Hz agreement shows two implementations of the same formula agree, and does not test the formula against physics. The audit nevertheless demonstrates the intended function of the protocol: applying the **(M4)** fields and the **(P3)** dependence test exposed within one step that the resource is synthetic, which is precisely the distinction those fields exist to reveal (Fig. 3).

3 Discussion

The protocol reframes what counts as a result in this literature. Under current practice a demonstration that a drone can be disturbed is publishable on its own; under the proposed protocol such a demonstration is a prompt to report **(M1)** through **(M5)**, without which it cannot be compared with anything. This raises the cost of an experiment and lowers the value of any single one, which we regard as the correct trade for a field whose central question is feasibility rather than existence.

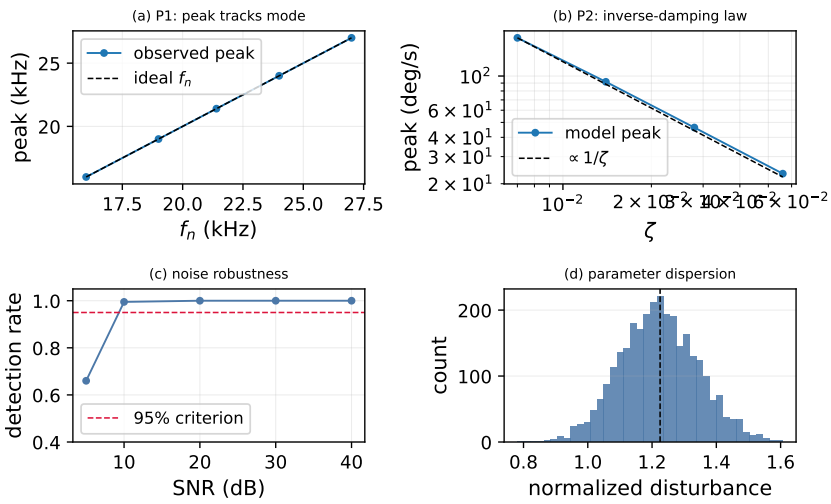


Fig. 2 Computational verification of the protocol. (a) **(P1)** peak tracking; (b) **(P2)** inverse-damping law on log-log axes; (c) mode-localization rate versus SNR, referenced to the resonant peak; (d) dispersion of the normalized disturbance under ±10% parameter uncertainty. Full numerical values are in Supplementary Information Tables S1–S5.

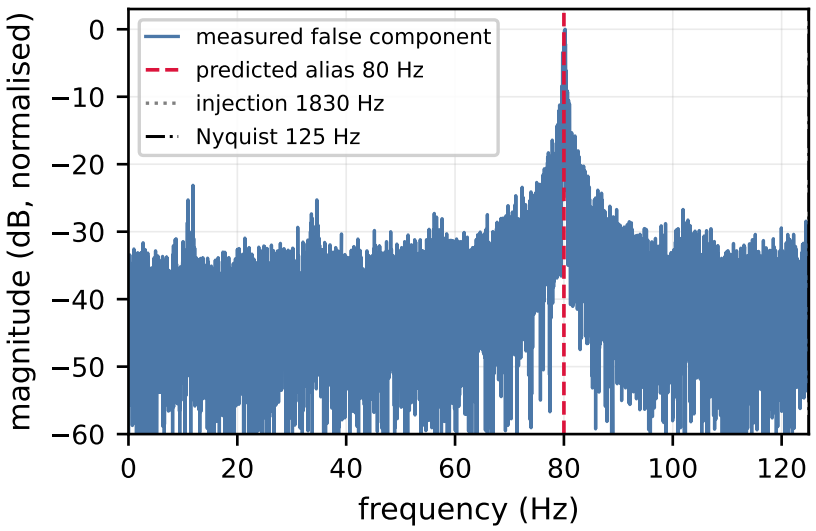


Fig. 3 Spectrum of the injected component in the IRIS_HITL resource, used for the provenance audit. The stated injection at 1830 Hz lies far above Nyquist while the dominant component sits at the predicted 80 Hz image. The resource performs software-only injection, so this is an audit of provenance and not a physical measurement.

The design also has a defensive payoff. Each model factor maps to a mitigation: ζ to damping, f_n to detuning, H_e and ε to notch and anti-alias filtering, f_s to randomized

sampling [4], and the tolerances to residual monitoring and physical challenge-response authentication [14]. Reporting the factors therefore tells a designer which mitigation applies, whereas reporting a disturbance level does not.

We considered and reject one alternative reading of our results: that because absolute index values depend on a poorly constrained κ_i, indices should be abandoned in favour of binary pass or fail tests. Binary tests are even less comparable, because a pass at an unreported pressure and incidence bounds nothing. The remedy is to report the factors, not to abandon quantification.

Several limitations bound what we claim. First, no physical acoustic measurement supports any claim in this paper; the only external data used is a software-injection resource whose implant we show to be synthetic. Second, the protocol has not been executed on hardware: its procedures are standard, but the precision requirements are derived from the model of Section 2.1 and may need revision once **(M2)** is attempted on real devices, where a calibrated free field at ultrasonic frequencies is not trivial to establish. Third, the five quantities constitute a basis, not a closed set: they are what the model identifies as necessary, and sufficiency is not established. Airframe acoustics, motor vibration masking and electromagnetic coupling lie outside the model, and **(P4)** is designed to detect the last of these by its signature insensitivity to incidence. Fourth, the tolerances in **(M5)** are stack-specific, so $\mathcal{V}$ is not a universal safety number but a device-and-stack-specific index comparable only when **(M5)** is reported with it. The assumptions on which the protocol rests, and the consequence of each being violated, are collected in Supplementary Table S8.

This paper proposed a measurement protocol for acoustic injection vulnerability of MEMS inertial sensors, derived five quantities with procedures, units and precisions, stated four falsifiable predictions with refutation conditions, and defined a minimum reportable unit. Its purpose is to make future studies comparable and feasibility claims testable. The comparability example shows why this matters: two existing studies cannot be ranked on what they publish, and the direction of any ranking supplied by assumption is not stable. The next step is execution, applying **(M1)** through **(M5)** to an instrumented IMU in a calibrated field, which would convert this specification into the first comparable dataset in the area.

4 Methods

This study involved no physical measurement, no human or animal subjects, and no new data collection. All results are computational. Equations (1)–(9) were implemented in Python (NumPy) on a frequency grid of 0.5–32 kHz at 1600 points, with saturation modelled as $\mathrm{sat}_Y(x) = Y \tanh(x/Y)$. The instantiated parameters, the per-prediction numerical results, the failure-mode tests, the distribution-sensitivity analysis and the external resource audit are given in full in Supplementary Information Sections S1–S6, and the three figures illustrating reporting format in Section S7. Stochastic steps are seeded and `experiments/run_all_experiments.py` regenerates every figure, table and result file in one command.

All data underlying this study are computationally generated and are archived with the submission as JSON result files (`results.json`, `results_computational.json`,

results_comparability.json, results_comparability_dist.json, results_empirical.json), provided together as the supplementary dataset file data_results.zip. All files are regenerated by a single command (experiments/run_all_experiments.py); all stochastic steps are seeded. The external IRIS_HITL files analysed in the provenance audit are publicly available from the UnRocker repository (Dataset/IRIS_HITL/) and are not redistributed here. Enquiries regarding data availability should be directed to Sen Fu, fs211314@gmail.com.

Data availability

All data underlying this study are computationally generated and are archived with the submission as JSON result files (results.json, results_computational.json, results_comparability.json, results_comparability_dist.json, results_empirical.json), provided together as the supplementary dataset file data_results.zip. All files are regenerated by a single command (experiments/run_all_experiments.py); all stochastic steps are seeded. The external IRIS_HITL files analysed in the provenance audit are publicly available from the UnRocker repository (Dataset/IRIS_HITL/) and are not redistributed here. Enquiries regarding data availability should be directed to Sen Fu, fs211314@gmail.com.

Code availability

All analysis scripts are included with the submission and are available from the corresponding author on request. run_all_experiments.py reproduces every figure, table and result file.

Author contributions

S.F. conceived the study, derived the reduced-order model, performed all computational analyses, generated the figures, and wrote the manuscript. S.T. contributed to the development of the measurement protocol and the falsifiable predictions, and revised the manuscript. Both authors reviewed and approved the final submitted version.

Funding

The authors received no funding for this work. No grants, contracts, or other financial support from any funding agency in the public, commercial, or not-for-profit sectors were obtained for this study.

Competing interests

The authors declare no competing interests.

Ethics statement

This study involved no human participants, no animal subjects and no new data collection. All results are computational.

References

[1] Son, Y., Shin, H., Kim, D., Park, Y., Noh, J., Choi, K., Choi, J., Kim, Y.: Rocking drones with intentional sound noise on gyroscopic sensors. In: 24th USENIX Security Symposium (USENIX Security 15), pp. 881–896 (2015). https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/son

[2] Trippel, T., Weisse, O., Xu, W., Honeyman, P., Fu, K.: WALNUT: Waging doubt on the integrity of MEMS accelerometers with acoustic injection attacks. In: 2017 IEEE European Symposium on Security and Privacy (EuroS&P), pp. 3–18 (2017). https://doi.org/10.1109/EuroSP.2017.42

[3] Khazaaleh, S., Korres, G., Eid, M., Rasras, M., Daqaq, M.F.: Vulnerability of MEMS gyroscopes to targeted acoustic attacks. IEEE Access **7**, 89534–89543 (2019) https://doi.org/10.1109/ACCESS.2019.2927084

[4] Jeong, J., Kim, D., Jang, J., Noh, J., Song, C., Kim, Y.: Un-rocking drones: Foundations of acoustic injection attacks and recovery thereof. In: Network and Distributed System Security Symposium (NDSS) (2023). https://doi.org/10.14722/ndss.2023.24112

[5] Yazdi, N., Ayazi, F., Najafi, K.: Micromachined inertial sensors. Proceedings of the IEEE **86**(8), 1640–1659 (1998) https://doi.org/10.1109/5.704269

[6] Cao, L.: Silicon based MEMS gyroscope structure and working principle. Dual-Mass Linear Vibration Silicon-Based MEMS Gyroscope, 13–32 (2023) https://doi.org/10.1007/978-981-19-9247-6_2

[7] Park, U., Lee, H., Han, H.: Response characteristics of a MEMS resonant accelerometer to external acoustic excitation. In: 2016 IEEE SENSORS, pp. 1–3 (2016). https://doi.org/10.1109/ICSENS.2016.7808717

[8] Manut, A., Syono, M.: Effects of mechanical geometries on resonance sensitivity of MEMS out-of-plane accelerometer. In: 2006 IEEE International Conference on Semiconductor Electronics (2006). https://doi.org/10.1109/SMELEC.2006.381013

[9] Zhang, Y., Zhang, W., Zhang, H.: Design and analysis of MEMS biaxial coupled resonance accelerometer. In: 2021 IEEE 16th International Conference on Nano/Micro Engineered and Molecular Systems (NEMS) (2021). https://doi.org/10.1109/NEMS51815.2021.9451525

[10] Olthuis, J.J., Arendsen, B., Sciancalepore, S., Zannone, N.: Impact of acoustic injection attacks on micro-UAVs. In: 2025 IEEE Conference on Communications and Network Security (CNS) (2025). https://doi.org/10.1109/CNS66487.2025.11195060

[11] Gao, M., Zhang, L., Shen, L., *et al.*: KITE: Exploring the practical threat from acoustic transduction attacks on inertial sensors. In: Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems (SenSys), pp. 696–709 (2022). https://doi.org/10.1145/3560905.3568532

[12] Sahu, C., Al-Hourani, A., Walia, S., Poddar, S.: RAiS dataset: Resonance-based Attacks on inertial Sensors. VectorNav VN100 IMU under acoustic interference with Xsens MTi-700 ground truth. Z-axis gyroscope resonance 27.55 kHz; 40 Hz sampling (2026). https://github.com/cmsahu97/RAiS-dataset

[13] Jeong, J., Kim, D., Jang, J., Noh, J., Song, C., Kim, Y.: UnRocker IRIS_HITL dataset: hardware-in-the-loop recordings of resonant acoustic injection on a PX4 flight controller. Contaminated channels attacc_accel0..2 paired with reference true_accel0..2; injection 1830 Hz at 250 Hz sampling (2023). https://github.com/jinseobjeong/UnRocker

[14] Shoukry, Y., Martin, P., Tabuada, P., Srivastava, M.: PyCRA: Physical challenge-response authentication for active sensors under spoofing attacks. Proceedings of the ACM Conference on Computer and Communications Security (2015) https://doi.org/10.1145/2810103.2813679

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- dataresults.zip