Geosci. Instrum. Method. Data Syst., 15, 195–211, 2026
https://doi.org/10.5194/gi-15-195-2026

Geoscientific Instrumentation Methods and Data Systems
Open Access
EGU

Causal and uncertainty-aware digital-twin framework for ultra–low-noise geoscientific inertial sensors

Antonino D'Alessandro

Istituto Nazionale di Geofisica e Vulcanologia, Osservatorio Nazionale, Rome, Italy

Correspondence: Antonino D'Alessandro (antonino.dalessandro@ingv.it)

Received: 2 February 2026 – Discussion started: 24 February 2026
Revised: 9 May 2026 – Accepted: 15 June 2026 – Published: 6 July 2026

Abstract. Ultra–low-noise inertial sensors are a cornerstone of modern geoscientific instrumentation, enabling high-resolution observations across seismology, geodesy, gravimetry, and vibration isolation. Achieving and reliably predicting their performance requires a rigorous treatment of physical causality, noise propagation, and uncertainty, particularly in force-feedback architectures operating near fundamental limits. In this study, we introduce a causal and uncertainty-aware digital-twin framework for the design and metrological assessment of ultra–low-noise geoscientific inertial sensors. The proposed framework integrates mechanical dynamics, force-feedback control, transduction, and digital acquisition within a physically realisable model that explicitly enforces causality and stability constraints. Starting from a minimal equation-of-motion description, the digital twin is formulated in the frequency domain to construct causal transfer functions and a comprehensive noise-budget model. The framework enables the systematic separation of fundamental thermal noise limits from implementation-dependent noise sources, including readout, actuation, and digital acquisition effects. We introduce quantitative performance metrics based on self-noise spectra, dominant noise regimes, crossover frequencies, and near-plateau bandwidths, allowing complex spectral behaviour to be condensed into actionable design indicators. Parameter uncertainties are propagated through the digital twin to provide uncertainty-aware performance estimates and robustness diagnostics. Through a series of illustrative analyses, we demonstrate how the proposed digital twin supports informed design trade-offs, identifies performance bottlenecks, and prevents non-physical or overly optimistic sensitivity estimates arising from non-causal modelling assumptions. While focused on inertial sensors, the methodology is general and transferable to other classes of geoscientific instruments. The framework provides a transparent and extensible foundation for next-generation sensor design, virtual experimentation, and metrologically consistent performance prediction.

1 Introduction

The design and metrological assessment of ultra-low-noise inertial sensors for geoscientific applications remains a central challenge in measurement science and engineering. These sensors – including broadband seismometers, compact accelerometers, and advanced gravimetric devices (Prasad et al., 2022) – are critical for resolving weak ground motions across broad frequency bands, enabling high-precision environmental monitoring, earthquake early warning, and advanced observational geodesy. In practice, instrument performance is constrained by a combination of fundamental thermomechanical limits and implementation-dependent noise sources such as thermal (Brownian) noise of mechanical suspensions, readout electronics noise, quantization error in digital acquisition, and control-loop artefacts (Ubhi et al., 2022; Ding et al., 2022; van Dongen et al., 2023).

Traditional sensor development has often proceeded through incremental improvements of individual components, with performance characterised post-facto via laboratory calibration and field testing. While such approaches are effective for benchmarking specific prototypes, they provide limited insight into the system-level trade-offs between sensitivity, bandwidth, dynamic range, and robustness, particularly when uncertainty propagation across the entire measurement chain is of interest. The need for rigorous uncer-

tainty frameworks in sensor design is well recognised in inertial navigation and metrology communities, where stochastic error modelling and error budgeting have been emphasised (El-Sheimy et al., 2020).

In parallel, the concept of a digital twin has emerged as a powerful paradigm for unifying physics-based simulation with real-time measurement data, enabling advanced prediction, optimisation, and uncertainty quantification in complex engineered systems. Digital twin methodologies have been extensively reviewed in industrial and infrastructure domains, highlighting their potential for model-based monitoring and decision support (Sun et al., 2025; Carlin et al., 2024). Recent developments have increasingly focused on uncertainty-aware and physics-informed digital twins for metrological applications, with particular emphasis on traceability, virtual experimentation, and hybrid modelling approaches that combine physical models with data-driven components (e.g., Wright and Davidson, 2024; Maculotti et al., 2024; Sun et al., 2024). In metrology, digital twins are gaining traction as tools for virtual calibration, uncertainty evaluation, and traceability analysis across diverse measurement systems, including coordinate measuring machines, optical scanners, and virtual metrological experiments (Vlaeyen et al., 2021; Poroskun et al., 2022).

Despite these advances, the digital-twin literature remains fragmented with respect to sensor-centric, causal, and uncertainty-aware frameworks that encompass both physical dynamics and digital acquisition processes. Many digital twin formulations in metrology still focus on high-level concepts or specific subsystems, without enforcing physical causality and realisability constraints essential for accurate modeling of closed-loop, force-feedback inertial sensors. Furthermore, the representation and propagation of uncertainty through the digital twin's internal variables and across the physical–digital boundary is often treated qualitatively or post hoc rather than systematically (Ríos et al., 2020).

In the field of inertial sensing, contemporary research continues to push performance boundaries through innovative mechanics and novel readout schemes, such as interferometric sensing with enhanced dynamic range and low noise floors (Kranzhoff et al., 2023), as well as advanced tilt and vibration measurement systems that delineate thermal and readout noise dominance across frequency (Bai et al., 2025). In particular, recent work has demonstrated significant progress in ultra-low-noise interferometric readout systems and compact inertial sensor architectures, pushing noise floors closer to fundamental limits while maintaining practical deployability (Carter et al., 2024; Carter et al., 2025; Bai et al., 2025). These efforts underscore the importance of integrated noise budgeting and uncertainty analysis, yet a comprehensive, unified framework that couples physical system dynamics, sensor signal chains, and digital data acquisition with rigorous uncertainty propagation is still absent.

This manuscript addresses these gaps by introducing a causal and uncertainty-aware digital-twin framework for ultra-low-noise geoscientific inertial sensors. The framework enforces physical causality and realisability constraints while systematically propagating uncertainty from mechanical parameters and control elements through to spectral performance metrics. It integrates mechanical modelling, transduction chains, force-feedback architectures, and digital acquisition within a single, physically consistent representation (Fig. 1). By decomposing self-noise into fundamental and implementation-dependent contributions and identifying dominant noise regimes, the framework facilitates quantitative evaluation of design choices and metrological limits.

Rather than relying on experimental datasets, the methodology is conceived as a design-stage tool to guide sensor architecture choices and parameter selection prior to hardware realisation. The remainder of this paper is organised as follows: Section 2 formalises design requirements and problem scope; Sect. 3 presents the digital-twin architecture; Sect. 4 develops the comprehensive noise model; Sects. 5 and 6 analyse spectral performance and design trade-offs; Sect. 7 synthesises performance metrics and optimisation strategies; Sect. 8 discusses broader implications and limitations; and Sect. 9 concludes with final remarks.

2 Design requirements and problem formulation

The design of ultra–low–noise inertial sensors for geoscientific applications constitutes a multi-objective optimisation problem in which sensitivity, bandwidth, dynamic range, robustness, and practical implementation constraints must be addressed simultaneously. Unlike post-deployment performance assessment or component-level optimisation approaches, the present study formulates the sensor design problem explicitly at the pre-implementation stage, adopting a system-level perspective aimed at predicting achievable performance prior to hardware realisation.

The conceptual scope of the problem is illustrated in Fig. 1, which summarises the functional elements of a force-feedback inertial sensor and their interactions. Within this framework, design requirements are not treated as independent specifications but as coupled constraints that jointly define the admissible design space.

The primary requirement considered in this work is the attainment of ultra–low self-noise levels over a broad frequency band spanning the ultra-low-frequency (ULF) regime to the classical broadband (BB) seismic band. Such performance is essential for a wide range of geoscientific applications, including broadband seismology, gravimetry, and long-term environmental monitoring, where weak ground motions must be resolved across several decades in frequency (Collette et al., 2012; D'Alessandro et al., 2019; Carter et al., 2024).

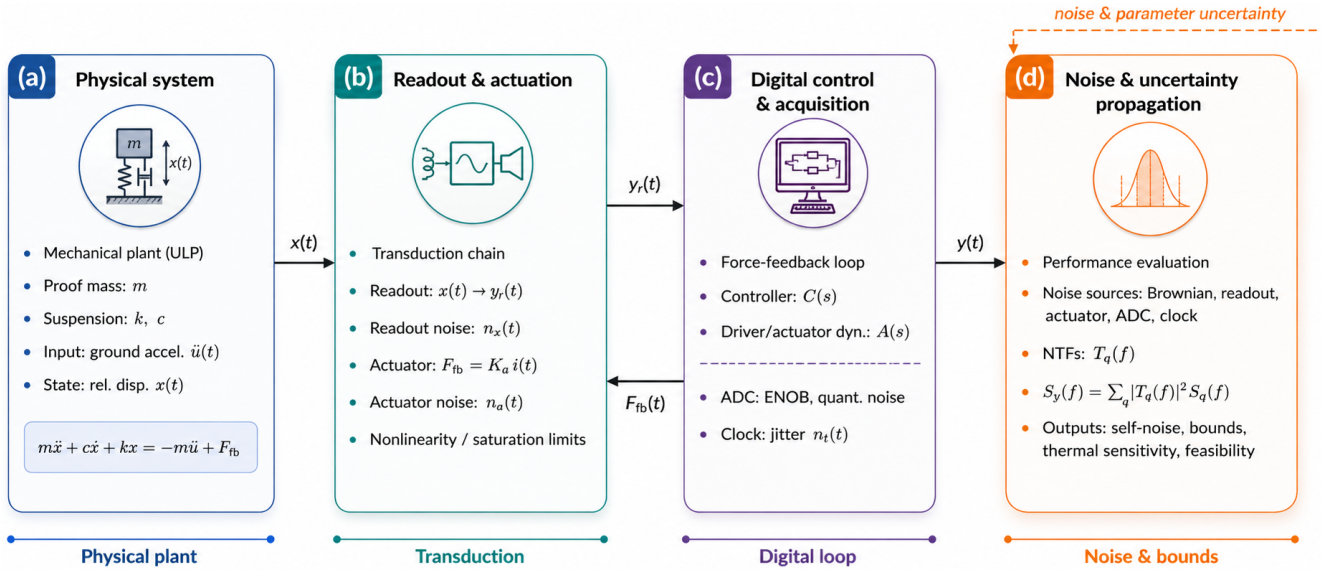

Figure 1. Conceptual block diagram of the digital-twin framework adopted for the analysis and design of a force-feedback ultra-long-period (ULP) seismometer. **(a)** Physical system: mechanical plant modeled as a proof mass–spring–damper system, driven by ground acceleration and controlled via feedback force, described by the equation of motion $m\ddot{x} + c\dot{x} + kx = -m\ddot{u} + F_{\mathrm{fb}}$. **(b)** Readout and actuation: transduction chain mapping the relative displacement $x(t)$ into the readout signal $y_{\mathrm{r}}(t)$, including readout noise and actuator dynamics, force generation, and nonlinearity or saturation effects. **(c)** Digital control and acquisition: closed-loop force-feedback architecture comprising the digital controller $C(s)$, driver/actuator dynamics $A(s)$, analog-to-digital conversion (ENOB and quantization noise), and clock jitter, yielding the output signal $y(t)$ in physical units. **(d)** Noise and uncertainty propagation: evaluation of the contribution of independent noise sources (Brownian, readout, actuator, ADC, and clock) through their corresponding noise transfer functions $T_q(f)$, resulting in the output power spectral density $S_y(f) = \sum_q |T_q(f)|^2 S_q(f)$. The framework enables quantitative assessment of self-noise, theoretical performance bounds, thermal sensitivity, and overall feasibility of the instrument design.

A second key requirement concerns dynamic range and saturation behaviour. Ultra-low-noise operation must coexist with the capability to accommodate large transient signals without loss of linearity, control stability, or actuator authority (Sider et al., 2023). In force-feedback architectures, this requirement couples mechanical design parameters, actuation strength, and control loop characteristics, precluding their independent optimisation (Ding et al., 2022).

Additional constraints arise from digital acquisition and system-level resources, including power consumption, telemetry bandwidth, and data resolution. Finite analog-to-digital converter (ADC) resolution, clock jitter, and digital signal processing limitations introduce noise contributions that may dominate the self-noise spectrum over specific frequency ranges, particularly at higher frequencies. These effects must therefore be incorporated explicitly into the design formulation rather than treated as secondary implementation details (El-Sheimy et al., 2020; Ubezio et al., 2023).

Finally, the design requirements implicitly assume a set of environmental and operational conditions, including mechanical coupling to the ground, thermal stability, and station infrastructure. While these factors are not modelled explicitly in the present framework, they define the operating context within which the sensor is expected to meet its performance objectives and delimit the interpretation of the resulting performance bounds.

On the basis of the above requirements, the design problem is formulated as follows: given a target sensitivity spectrum, a frequency band of interest, and a set of implementation constraints, determine whether a physically realisable inertial sensor architecture can achieve the desired performance and identify the dominant mechanisms limiting that performance.

A central aspect of this formulation is the explicit separation between fundamental physical limits, such as thermal noise associated with mechanical dissipation, and implementation-dependent limits arising from readout electronics, actuation mechanisms, feedback control, and digital acquisition. This distinction is essential for meaningful interpretation of noise budgets and for guiding design decisions toward genuinely performance-limiting components rather than secondary contributors (Collette et al., 2012; Maculotti et al., 2024).

The sensor is modelled as a linear, time-invariant system operating around a stable equilibrium point. Its behaviour is described through transfer functions linking ground acceleration to the measured output quantity. Nonlinear effects such as actuator saturation, hysteresis, and large-amplitude geometric nonlinearities are neglected. These assumptions

are standard in the analysis of ultra-low-noise inertial sensors and are justified when the focus is on noise-limited performance under nominal operating conditions (Carter et al., 2024).

Within this framework, the observable of interest is the output acceleration estimate, whose power spectral density is determined by the combined effect of the sensor dynamics and all internal noise sources propagated through the system. The design objective is therefore the optimisation of the total self-noise spectrum, subject to causality, stability, and realisability constraints, rather than the minimisation of individual noise sources in isolation.

The formulation adopted here deliberately excludes experimental calibration data and site-specific noise conditions. This choice reflects the intended role of the framework as a design-stage and feasibility-assessment tool, rather than as a post-deployment diagnostic method. While this limits direct comparison with specific instruments, it enables the identification of general trends, trade-offs, and performance bounds that are transferable across sensor classes and deployment scenarios (Wright and Davidson, 2024; Ríos et al., 2020).

Model parameters, including mechanical damping and electronic noise levels, are assumed to be characterised by probability distributions reflecting their uncertainty. The systematic propagation of these uncertainties through the digital-twin framework is addressed in subsequent sections. The implications of these assumptions, and their potential impact on real-world performance, are discussed in Sect. 8.

3 Digital-twin architecture and causal modeling

The digital twin developed in this study is conceived as a physics-based, causal, and uncertainty-aware representation of an ultra–low-noise inertial sensor, explicitly designed to support metrological analysis and design-stage optimisation. Unlike generic simulation models or reduced-order representations, the proposed digital twin integrates mechanical dynamics, transduction, actuation, control, and digital acquisition within a unified framework that preserves physical causality and realisability.

A frequency-domain representation of the digital-twin response is shown in Fig. 2, illustrating the effect of force-feedback control on the system dynamics. This architecture reflects the functional decomposition of a force-feedback inertial sensor and provides the basis for subsequent noise propagation and performance analysis. In particular, Fig. 2a compares the open-loop mechanical response with the closed-loop response obtained under force-feedback control. The open-loop configuration exhibits a pronounced resonance at the natural frequency, whereas the closed-loop system suppresses this resonance and reduces proof-mass motion over a broad frequency band. Figure 2b shows the corresponding force-balance transfer function, highlighting how the feedback force tracks the inertial force at low frequencies and progressively rolls off at higher frequencies due to the finite bandwidth of the control loop.

The mechanical subsystem is modelled as a single-degree-of-freedom inertial plant characterised by an effective mass, elastic stiffness, and dissipative damping. This representation captures the dominant dynamics governing the sensor response in the frequency range of interest and is commonly adopted in the analysis of broadband and ultra-low-noise inertial sensors (Collette et al., 2012; D'Alessandro et al., 2019; Carter et al., 2024).

The validity of the single-degree-of-freedom (SDOF) approximation is restricted to the frequency range in which the fundamental mode dominates the mechanical response. In practice, this corresponds to frequencies sufficiently below the first higher-order structural resonance of the sensor assembly, where multi-mode effects and internal deformation of the mechanical structure can be neglected (Collette et al., 2012).

At very low frequencies, additional effects such as tilt coupling, foundation compliance, and environmental interactions may introduce deviations from the idealised inertial response. These effects are not explicitly modelled in the present framework and therefore delimit the lower bound of applicability of the SDOF representation.

The model further assumes rigid coupling between the sensor frame and the ground, negligible rotational degrees of freedom, and operation in the linear regime around a stable equilibrium configuration. Under these conditions, the SDOF approximation provides an accurate and physically consistent description of the dominant sensor dynamics relevant for noise-limited performance analysis.

Thermal noise associated with mechanical dissipation is treated as an intrinsic property of the plant and constitutes a fundamental performance limit. The mechanical model is assumed to operate in the linear regime around a stable equilibrium point, and temperature is treated as a stationary parameter. Spatially distributed modes, geometric nonlinearities, and thermoelastic coupling effects are neglected, an assumption justified when focusing on noise-limited performance well below structural resonance frequencies.

The transduction stage converts the mechanical state of the plant into an electrical signal. The digital twin represents this stage through a linear gain and an additive readout noise term, which may encompass optical, capacitive, or electromagnetic sensing mechanisms depending on the sensor implementation. This abstraction allows the framework to remain agnostic with respect to the specific readout technology while retaining its metrological relevance. Compact interferometric readout implementations have recently been integrated in suspended interferometers, providing practical benchmarks for low-noise readout assumptions (Carter et al., 2025; Mitchell et al., 2025).

Readout noise is explicitly modelled as a stochastic process whose spectral characteristics can be prescribed or parameterised based on design assumptions. This representa-

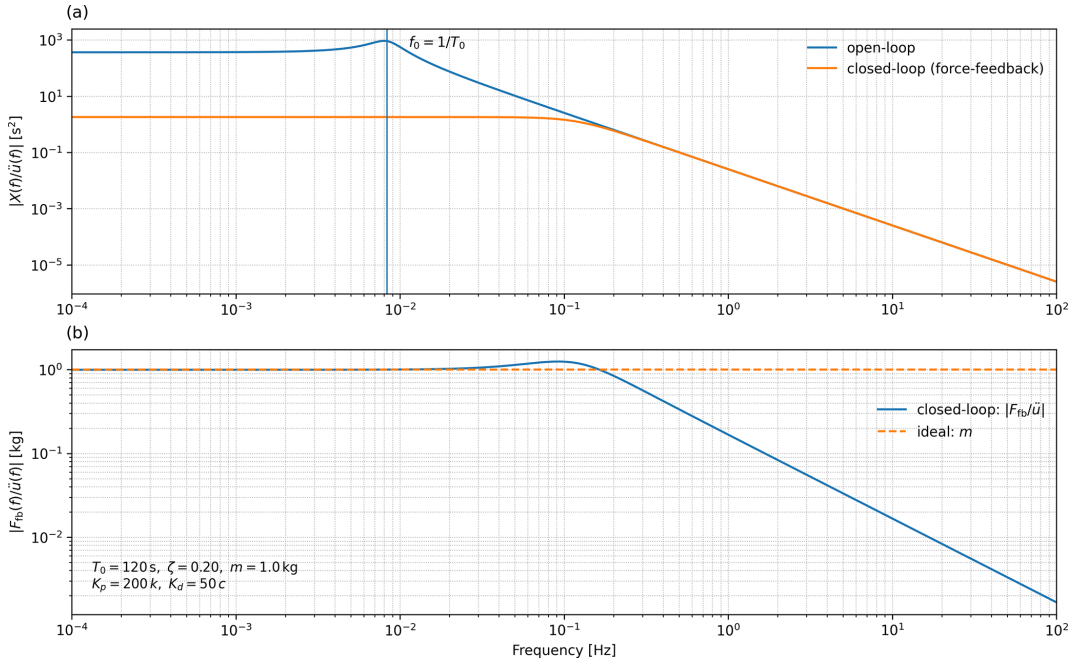

Figure 2. Frequency-domain response of the ultra-long-period (ULP) force-feedback seismometer digital twin. **(a)** Magnitude of the mechanical compliance $|X(f)/\ddot{u}(f)|$, comparing the open-loop response of the mechanical plant with the closed-loop response obtained through force-feedback control. The open-loop resonance at the natural frequency $f_0 = 1/T_0$ is clearly visible, while the closed-loop configuration suppresses the resonance and significantly reduces proof-mass motion over a broad frequency band. **(b)** Magnitude of the force-balance transfer function $|F_{\mathrm{fb}}(f)/\ddot{u}(f)|$ in the closed-loop configuration, compared with the ideal low-frequency limit $F_{\mathrm{fb}}/\ddot{u} = m$. At low frequencies, the feedback force accurately tracks the inertial force associated with ground acceleration, whereas at higher frequencies the response rolls off due to the finite bandwidth of the control loop. The parameters reported in the inset correspond to a representative ULP configuration and a proportional–derivative feedback law.

tion is consistent with recent advances in ultra-low-noise front-end electronics for inertial sensors, where circuit-level optimisation plays a critical role in approaching fundamental sensitivity limits (Chen et al., 2024). This approach is consistent with modern inertial sensor analyses, where readout noise often dominates the self-noise spectrum outside the thermal-noise-limited band (Ding et al., 2022; Carter et al., 2024).

Force-feedback actuation is incorporated into the digital twin to stabilise the mechanical plant and to linearise the sensor response. The actuation subsystem is represented by a causal transfer function linking the control signal to an applied force on the mechanical mass. Actuator noise and finite authority are treated as implementation-dependent limitations and are included explicitly in the model.

The control architecture is described by a stabilising controller acting on the measured output. Particular care is taken to ensure that the closed-loop system satisfies causality and stability constraints, as these conditions directly influence noise propagation and achievable performance. Non-causal inversions or idealised feedback laws are deliberately avoided, in contrast to some analytical treatments that neglect realisability constraints, often relying on idealised or non-causal representations of system transfer functions (Oppenheim and Schafer, 2010; Bendat and Piersol, 2010).

The final stage of the digital twin represents the digital acquisition process, including analog-to-digital conversion, finite resolution, and timing uncertainty. Quantisation noise and clock jitter are modelled explicitly and propagated through the system in subsequent analyses. This treatment reflects the growing recognition that digital acquisition can impose dominant performance limits in ultra-low-noise sensors, particularly at higher frequencies or under stringent power constraints (El-Sheimy et al., 2020; Ubezio et al., 2023).

Discretisation effects are treated consistently with the causal structure of the system. Continuous-time transfer functions are mapped to their discrete-time counterparts using stable and physically meaningful transformations. This ensures that the digital twin remains a faithful representation of a realisable sensor system rather than an idealised mathematical construct. Open-source simulation toolchains widely used in precision interferometry provide a reproducibility baseline for digital-twin implementations (Brown et al., 2020).

All model parameters within the digital twin are assumed to be affected by uncertainty arising from manufacturing tolerances, environmental variability, and modelling approxi-

mations. These uncertainties are represented through probability distributions assigned to the relevant parameters and are propagated through the digital twin in subsequent sections.

It is assumed that parameter uncertainties are statistically independent unless otherwise stated. While this assumption may not hold in all practical cases, it provides a tractable starting point for uncertainty propagation and is consistent with existing metrological digital-twin frameworks (Wright and Davidson, 2024; Maculotti et al., 2024). The implications of this assumption are discussed in Sect. 8.

Within the proposed digital-twin framework, the mechanical plant and the force-feedback actuation are compactly described by the equation of motion

$$m\ddot{x}(t) + c\dot{x}(t) + kx(t) = -ma_{\mathrm{g}}(t) + F_{\mathrm{fb}}(t), \tag{1}$$

where $x(t)$ denotes the relative displacement of the inertial mass with respect to the sensor frame, $a_{\mathrm{g}}(t)$ is the ground acceleration to be measured (in Fig. 1, the same quantity is represented using the kinematic notation $\ddot{u}$), m, c, and k represent the effective mass, damping coefficient, and stiffness of the mechanical plant, and $F_{\mathrm{fb}}(t)$ is the feedback force applied by the actuator. Equation (1) provides the physical backbone of the digital twin, from which the causal transfer functions and noise propagation relationships are derived in the frequency domain in the following section.

4 Metrological noise modeling and uncertainty propagation

The metrological performance of an ultra–low-noise inertial sensor is ultimately determined by the propagation of multiple stochastic noise sources through the causal dynamics of the sensor system. Building on the physical model introduced in Sect. 3, this section formalises the noise modeling framework adopted in the digital twin and defines the methodology used to propagate uncertainty from individual noise sources to the total self-noise spectrum.

A schematic overview of the noise budget and its decomposition into individual contributions is shown in Fig. 3, which provides a graphical representation of the formalism developed below.

Starting from the equation of motion introduced in Eq. (1), the system is transformed into the frequency domain under the assumption of linear, time-invariant dynamics. Taking the Fourier transform of Eq. (1) and using the standard correspondence $\mathrm{d}/\mathrm{d}t \rightarrow \mathrm{i}\omega$, the equation of motion can be written in the frequency domain as

$$(-m\omega^2 + \mathrm{i}c\omega + k)X(\omega) = -mA_{\mathrm{g}}(\omega) + F_{\mathrm{fb}}(\omega),$$

Denoting Fourier-transformed quantities by capital letters, the relative displacement $X(\omega)$ of the inertial mass can be written as

$$X(\omega) = \frac{-mA_{\mathrm{g}}(\omega) + F_{\mathrm{fb}}(\omega)}{k - m\omega^2 + \mathrm{i}c\omega}, \tag{2}$$

where $A_{\mathrm{g}}(\omega)$ is the ground acceleration spectrum and $F_{\mathrm{fb}}(\omega)$ represents the feedback force in the frequency domain. Equation (2) defines the mechanical susceptibility of the plant and provides the basis for constructing the transfer functions linking each noise source to the sensor output.

In a force-feedback configuration, the measured output is typically proportional to an estimate of the ground acceleration derived from the control signal and the plant response. The digital twin represents this relationship through a set of causal transfer functions $T_q(\omega)$, each associated with a specific noise source q.

The total self-noise of the sensor is defined as the output noise spectrum obtained in the absence of ground motion, i.e. for $A_{\mathrm{g}}(\omega) = 0$. Under the assumption of linearity and mutual incoherence of the noise sources, the total output power spectral density can be expressed as the sum of the individual contributions propagated through their respective transfer functions. This follows from standard results in linear systems theory, where each noise source is treated as an independent stochastic input and its contribution is weighted by the squared magnitude of the corresponding transfer function (e.g., Bendat and Piersol, 2010; El-Sheimy et al., 2020). Under this condition, the output acceleration noise power spectral density (PSD) can be expressed as the incoherent sum of the contributions from all internal noise sources,

$$S_{\mathrm{a,self}}(\omega) = \sum_q |T_q(\omega)|^2 S_q(\omega), \tag{3}$$

where $S_q(\omega)$ denotes the PSD of the qth noise source and $T_q(\omega)$ is the corresponding transfer function from that source to the output acceleration estimate. Equation (3) constitutes the central metrological relation of the framework and formalises the noise-budget decomposition illustrated in Fig. 3.

The noise sources considered in this work include: (i) thermal (Brownian) noise associated with mechanical dissipation, (ii) readout noise originating from the transduction stage, (iii) actuation noise introduced by the force-feedback mechanism, (iv) digital acquisition noise, including quantisation and timing uncertainty.

Thermal noise associated with mechanical damping represents a fundamental physical limit that cannot be reduced without modifying the underlying dissipation mechanisms. Its contribution is governed by the fluctuation–dissipation theorem and depends on temperature, damping, and mechanical susceptibility (Collette et al., 2012).

In contrast, readout, actuation, and digital acquisition noises are implementation-dependent and reflect technological and design choices. Their relative importance varies across frequency and design parameter space, and they often dominate the self-noise spectrum outside the thermal-noise-limited band (Ding et al., 2022; Carter et al., 2024). The ex-

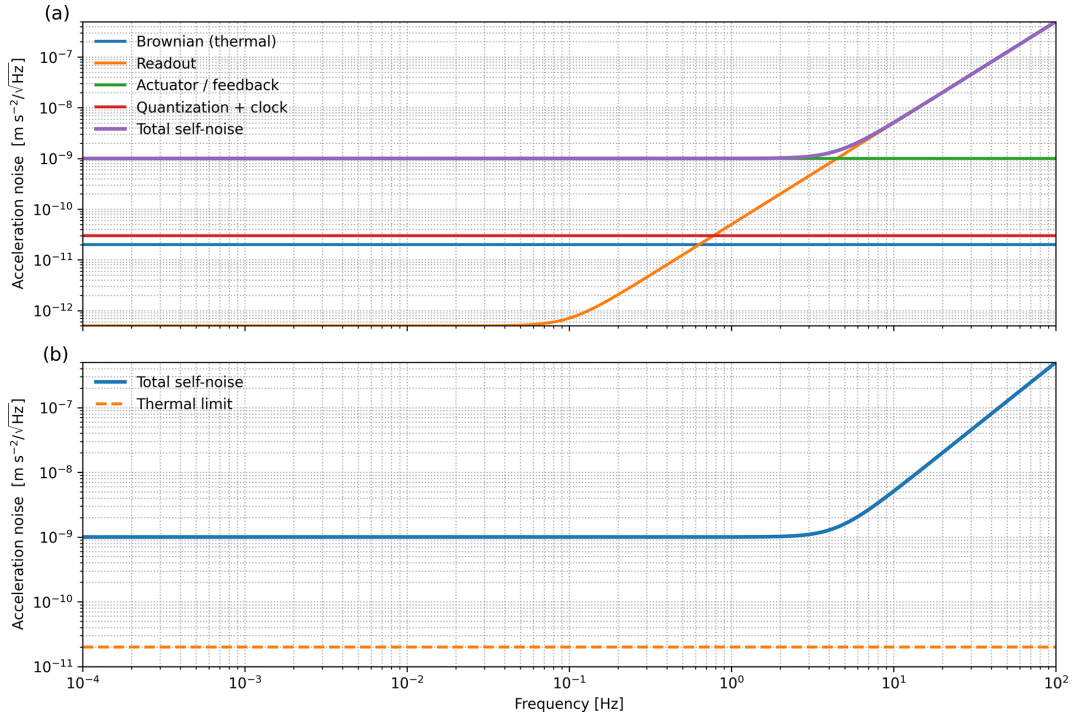

Figure 3. Frequency-dependent acceleration self-noise budget of the force-feedback ULP seismometer derived from the digital-twin framework. **(a)** Individual noise contributions referred to input acceleration, including Brownian (thermal) noise of the mechanical plant, readout noise, actuator/feedback noise, and quantization plus clock jitter, together with their quadratic sum defining the total self-noise. The spectrum highlights the transition from readout- and quantization-limited performance at low frequencies to feedback- and control-limited behavior at higher frequencies. **(b)** Comparison between the total self-noise and the fundamental thermal limit set by the mechanical suspension. Model parameters are indicated in the panel and correspond to a nominal ultra–long-period design with force-feedback control. The separation between the total self-noise and the thermal limit quantifies the residual performance gap attributable to non-thermal noise sources and control-loop implementation.

plicit separation between these two classes of noise sources is essential for interpreting metrological bounds and guiding effective design optimisation.

All noise sources and model parameters entering Eq. (3) are subject to uncertainty arising from manufacturing tolerances, environmental variability, and modelling assumptions. Within the digital-twin framework, these uncertainties are represented through probability distributions assigned to the relevant parameters.

Uncertainty propagation is performed by evaluating the statistical distribution of the self-noise spectrum resulting from Eq. (3), given the distributions of $S_q(\omega)$ and the parameters defining $T_q(\omega)$. This approach is consistent with recent developments in digital twin-based uncertainty mapping and error propagation in metrological systems, where spatial and parametric uncertainties are explicitly quantified within virtual representations (Sepahi-Boroujeni and Khameneifar, 2024). This approach enables the estimation of confidence intervals on the predicted self-noise and provides a quantitative measure of robustness with respect to parameter variability, in line with modern metrological digital-twin concepts (Scholz et al., 2022; Wübbeler et al., 2022; Wright and Davidson, 2024; Maculotti et al., 2024).

Parameter correlations are neglected unless explicitly stated. While this assumption may not hold in all practical implementations, it allows the identification of dominant uncertainty drivers and represents a reasonable first-order approximation for design-stage analysis.

The noise modeling framework assumes linearity, stationarity, and mutual incoherence of the noise sources. Cross-correlations between noise processes and non-stationary effects are not considered. These assumptions are standard in inertial sensor metrology and are justified when the objective is to characterise noise-limited performance under nominal operating conditions.

The implications of these assumptions, and their potential impact on real-world sensor behaviour, are discussed in Sect. 8.

5 Spectral performance and self-noise regimes

This section analyses the spectral performance predicted by the digital twin and characterises the resulting self-noise in

terms of dominant noise regimes across frequency. The objective is to translate the metrological formulation introduced in Sect. 4 into physically interpretable performance metrics that are directly relevant for sensor design.

A representative example of the self-noise spectrum and its decomposition into individual contributions is shown in Fig. 4, which illustrates how fundamental and implementation-dependent noise sources shape the achievable performance across the frequency band of interest.

The parameter ranges adopted in this analysis, including mechanical properties, damping ratios, and representative noise levels, are consistent with those reported for state-of-the-art broadband and ultra–low-noise inertial sensors (e.g., Collette et al., 2012; Carter et al., 2024), ensuring that the simulated performance remains representative of physically achievable instrument designs.

The total self-noise spectrum $S_{\mathrm{a,self}}(\omega)$, defined in Eq. (3), typically exhibits distinct frequency-dependent behaviours governed by the interplay between mechanical susceptibility and noise transfer functions. Of particular interest is the presence of noise plateaus, i.e. frequency intervals over which the self-noise remains approximately constant.

To formalise this concept, the acceleration self-noise amplitude spectral density (ASD) is defined as

$$N_{\mathrm{a}}(\omega) = \sqrt{S_{\mathrm{a,self}}(\omega)}. \tag{4}$$

A plateau region is identified when the logarithmic slope of $N_{\mathrm{a}}(\omega)$ with respect to frequency satisfies

$$\left| \frac{\mathrm{d}\log N_{\mathrm{a}}(\omega)}{\mathrm{d}\log \omega} \right| < \varepsilon, \tag{5}$$

where ε is a small threshold chosen to discriminate between flat and sloped spectral behaviour. While the precise value of ε is application-dependent, the qualitative identification of plateau regions is robust with respect to reasonable threshold variations.

Plateaus are of particular metrological relevance because they define frequency bands in which the sensor sensitivity is maximised and least sensitive to modelling uncertainties.

Across the full frequency band, different noise sources dominate the self-noise spectrum. At low frequencies, the response is typically governed by thermal noise associated with mechanical damping, reflecting the fundamental limit imposed by dissipation mechanisms. In intermediate frequency ranges, readout or actuation noise may dominate, depending on the chosen transduction and control architecture. At higher frequencies, digital acquisition noise, including quantisation and timing uncertainty, often becomes the limiting factor.

These dominant noise regimes are identified by comparing the individual terms in Eq. (3) and determining, at each frequency, the noise source contributing the largest fraction to $S_{\mathrm{a,self}}(\omega)$. This regime-based interpretation is illustrated in Fig. 4, where transitions between thermal-limited and implementation-limited behaviour are clearly visible.

The identification of dominant regimes is a consolidated result, as it directly follows from the noise budget formalism and does not depend on subjective interpretation. In contrast, the precise frequency boundaries between regimes should be regarded as model-dependent and may shift as design parameters or uncertainty assumptions are varied.

A central outcome of the spectral analysis is the explicit comparison between the total self-noise and the thermal noise floor. The ratio

$$R(\omega) = \frac{S_{\mathrm{a,self}}(\omega)}{S_{\mathrm{a,th}}(\omega)} \tag{6}$$

provides a frequency-dependent measure of the margin to the thermal limit, where $S_{\mathrm{a,th}}(\omega)$ denotes the thermal noise contribution. Values of $R(\omega)$ close to unity indicate near-thermal-limited performance, whereas larger values highlight the dominance of implementation-dependent noise sources.

This ratio constitutes a key diagnostic metric for sensor design, as it directly indicates whether further performance improvements require fundamental changes to the mechanical plant or, alternatively, technological improvements in readout, actuation, or digital acquisition. The interpretation of $R(\omega)$ across frequency is illustrated in Fig. 4, where near-thermal and implementation-limited bands can be clearly distinguished.

The spectral features identified in this section are a direct consequence of the causal structure of the digital twin and the noise modeling assumptions introduced in Sect. 4. While the existence of plateaus and dominant noise regimes is a robust qualitative result, their quantitative characteristics depend on model parameters and uncertainty distributions.

Nonlinear effects, cross-correlations between noise sources, and non-stationary behaviour are not considered in this analysis. These factors may alter the detailed spectral structure in real instruments, particularly under extreme operating conditions. Nevertheless, the regime-based interpretation presented here provides a physically meaningful and practically useful framework for guiding sensor design.

6 Crossover frequencies and design trade-offs

While the identification of dominant noise regimes provides a qualitative understanding of sensor performance, practical design decisions require quantitative metrics capable of capturing transitions between regimes and their dependence on key design parameters. In this section, such metrics are introduced through the concept of crossover frequencies, which mark the boundaries between noise-dominated regimes and provide a compact description of performance trade-offs.

A representative mapping of crossover frequencies and dominant regimes in the design parameter space is shown in Fig. 6, which synthesises the spectral analyses discussed in the previous section.

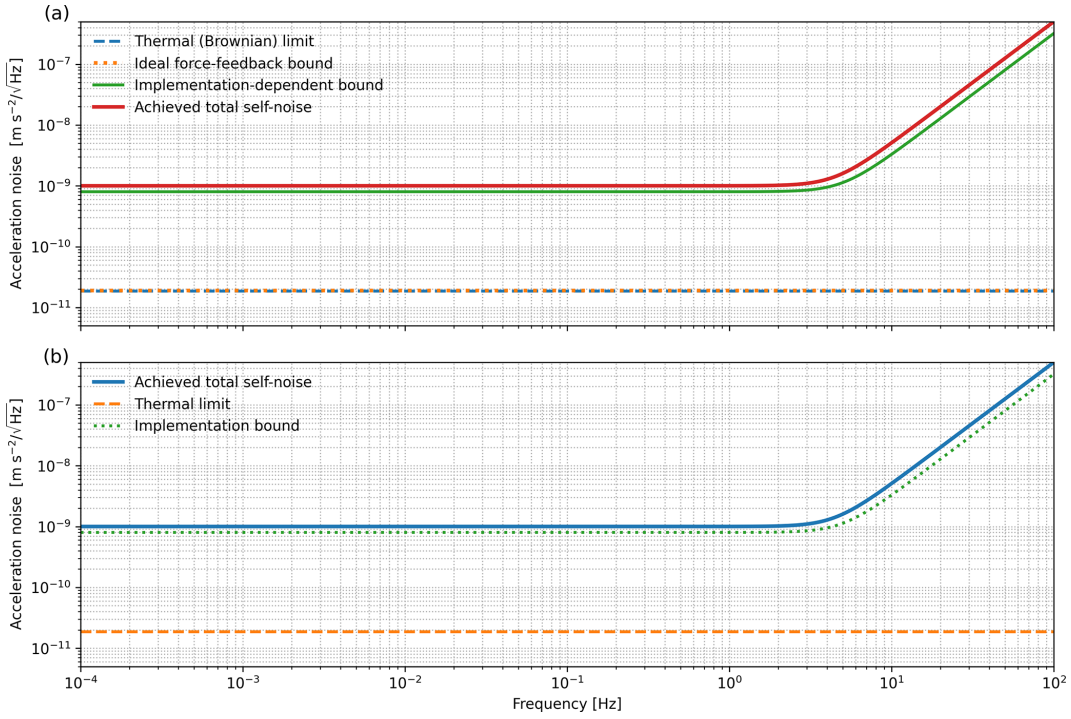

Figure 4. Metrological acceleration-noise bounds and achieved performance of the force-feedback ultra–long-period (ULP) seismometer derived from the digital-twin framework. **(a)** Fundamental and implementation-dependent noise bounds referred to input acceleration. The thermal (Brownian) limit represents the irreducible noise floor imposed by mechanical dissipation in the suspension. The ideal force-feedback bound coincides with the thermal limit in the absence of electronic and control noise and is shown with a slight offset for visual clarity. The implementation-dependent bound accounts for practical non-idealities introduced by actuation, digitization, timing, and control-loop dynamics, and exhibits a frequency-dependent rise imposed by closed-loop causality. The achieved total self-noise is shown for reference. **(b)** Comparison between the achieved total self-noise and the theoretical bounds. The separation from the thermal limit quantifies the residual performance gap attributable to control-loop implementation and electronic noise sources, providing a direct metric to assess design margins and guide further optimization toward fundamental physical limits.

For any pair of noise sources q_1 and q_2, a crossover frequency $\omega_c^{(q_1,q_2)}$ is defined as the solution of

$$\left|T_{q_1}(\omega_c)\right|^2 S_{q_1}(\omega_c) = \left|T_{q_2}(\omega_c)\right|^2 S_{q_2}(\omega_c), \tag{7}$$

i.e. the frequency at which the contributions of the two noise sources to the total self-noise spectrum are equal. Below and above this frequency, the dominant contribution switches from one noise source to the other.

In practice, the most relevant crossover frequencies involve transitions between thermal noise and implementation-dependent noise sources, such as readout or digital acquisition noise. These transitions delineate frequency intervals in which further performance improvements require fundamentally different design strategies.

Building on the plateau definition introduced in Sect. 5, the near-plateau bandwidth is defined as the frequency interval over which the self-noise remains within a prescribed margin of the minimum achievable level. Formally, this interval is given by

$$N_a(\omega) \le (1+\delta) N_{a,\min}, \tag{8}$$

where $N_{a,\min}$ is the minimum value of the self-noise ASD and δ is a tolerance parameter. The near-plateau bandwidth provides a concise metric for assessing the usable frequency range over which the sensor operates close to its optimal sensitivity.

As illustrated in Fig. 6, this bandwidth is strongly dependent on implementation parameters such as ADC resolution and control-loop design, and it often represents a more informative performance metric than the absolute minimum noise level alone.

The crossover frequencies and near-plateau bandwidth jointly define a design trade-off space in which improvements along one dimension may degrade performance along another. For example, increasing digital resolution can extend the near-plateau bandwidth toward higher frequencies but may impose penalties in power consumption or system complexity. Conversely, modifying mechanical damping may shift thermal-to-readout crossover frequencies at the expense of increased sensitivity to environmental perturbations.

The regime maps shown in Fig. 6 summarise these trade-offs by identifying, for each region of the parameter space,

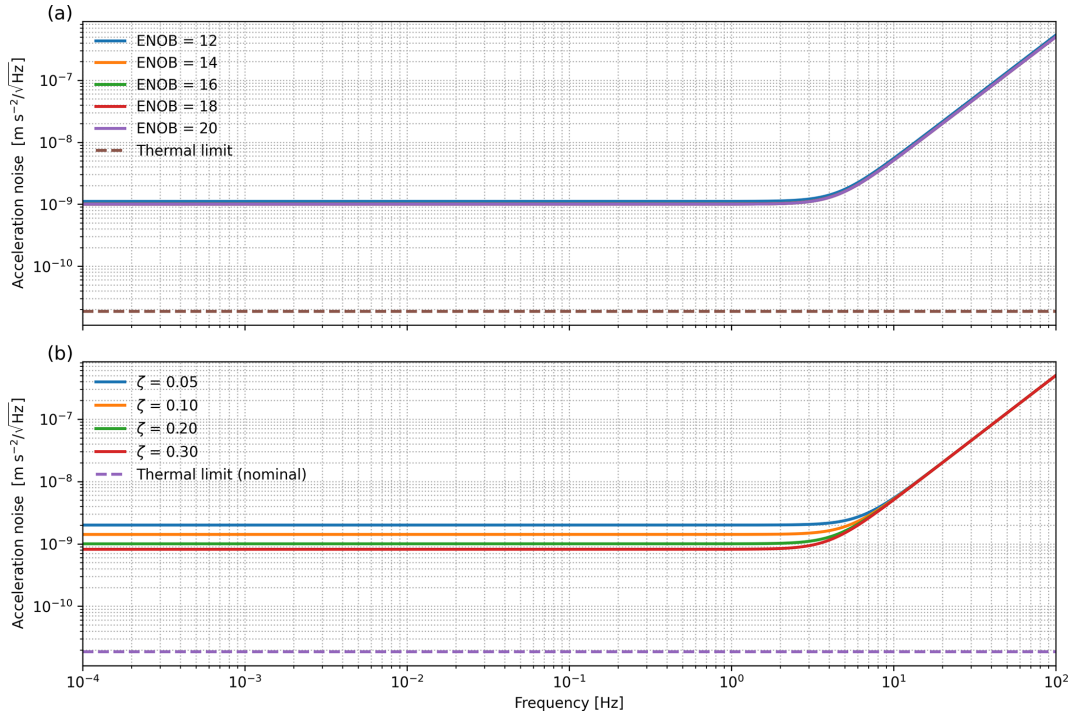

Figure 5. Spectral sensitivity of the closed-loop instrument self-noise to key design parameters. **(a)** Total acceleration self-noise amplitude spectral density $S_a^{1/2}(f)$ for different effective ADC resolutions (ENOB = 12–20 bits), computed for a fixed mechanical plant ($m = 1\,\text{kg}$, $T_0 = 120\,\text{s}$, $\zeta = 0.20$) and identical control architecture. At low frequencies the spectra collapse onto a common plateau set by non-digital noise sources, whereas at higher frequencies the achievable noise floor progressively improves with increasing ENOB as quantization and digital-chain contributions are reduced. The dashed curve indicates the Brownian (thermal) acceleration limit, shown for reference. **(b)** Sensitivity of the total acceleration self-noise to the damping ratio ζ (0.05–0.30) at fixed ENOB (16 bits). Variations in ζ primarily affect the low-frequency noise level through the balance between thermal dissipation and feedback-related contributions, while the high-frequency rise remains controlled by the digital and readout bandwidth. The dashed line marks the nominal thermal limit. Together, the two panels illustrate how electronic resolution and mechanical damping act on distinct frequency regimes, defining complementary design trade-offs for ultra-low-frequency force-feedback inertial sensors.

the noise source that ultimately limits performance. These maps constitute a consolidated result, as they are derived directly from the noise-budget formalism and the causal structure of the digital twin. However, the precise boundaries between regimes should be interpreted as model-dependent, reflecting assumptions on parameter uncertainty and noise spectra.

From a design perspective, crossover-based metrics provide actionable guidance by indicating whether further optimisation efforts should focus on fundamental mechanical improvements or on technological enhancements in readout, actuation, or digital acquisition. In this sense, crossover frequencies act as decision thresholds separating regimes where different design strategies are effective.

It is important to note that crossover frequencies are not intrinsic properties of the sensor but emerge from the interaction between physical dynamics and implementation choices. As such, they should be interpreted within the context of the assumed operating conditions and uncertainty model. The broader implications of this dependence are discussed in Sect. 8.

7 Performance metrics and design optimization

The spectral analyses and crossover-based diagnostics introduced in the previous sections provide detailed insight into the noise-limited behaviour of ultra–low-noise inertial sensors. For design purposes, however, it is often desirable to condense this information into a limited set of performance metrics that can guide optimisation decisions and enable comparisons between alternative architectures. This section introduces such metrics and illustrates their role within the digital-twin framework.

A synthesis of the performance metrics discussed below and their dependence on key design parameters is shown in Fig. 7, which summarises the design space explored in this study.

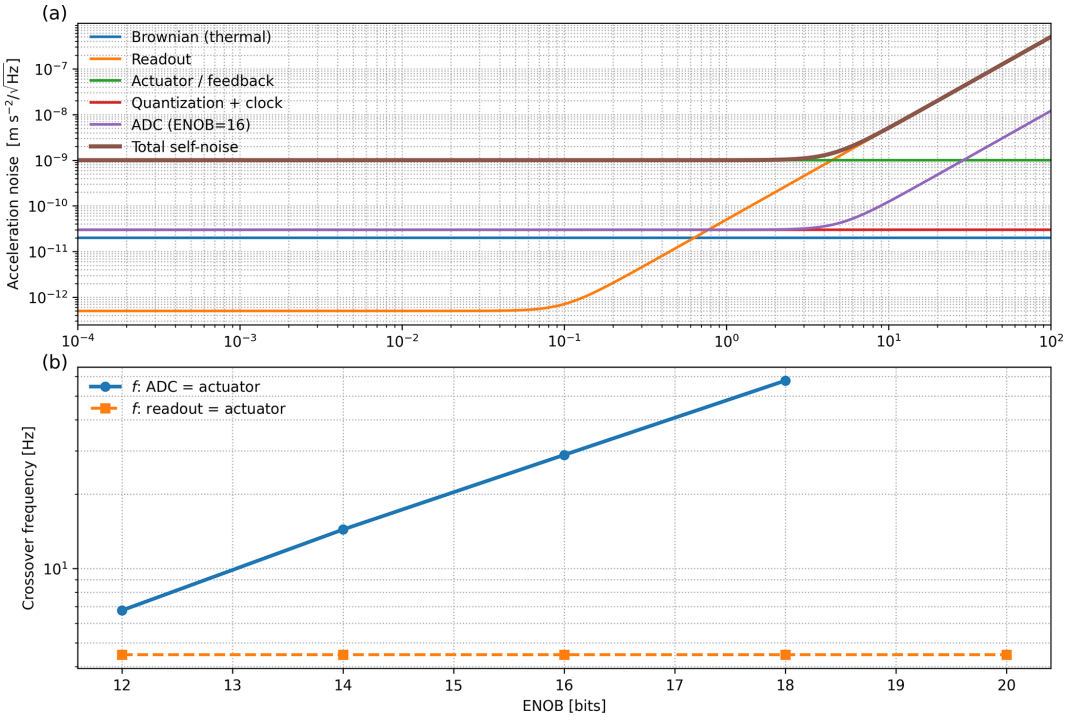

Figure 6. Noise regimes and crossover frequencies in a digitally controlled broadband inertial sensor. **(a)** Acceleration noise amplitude spectral density (ASD) as a function of frequency, showing the individual contributions from Brownian (thermal) noise, readout noise, actuator/feedback noise, quantization and clock noise, and the digital (ADC) contribution for a nominal resolution of ENOB = 16 bits. The resulting total self-noise (quadratic sum of all components) is also shown. The plot highlights the frequency-dependent dominance of different noise sources, with readout noise governing the low-frequency range, actuator/feedback noise setting the mid-band floor, and digital/ADC-related effects becoming relevant at higher frequencies due to closed-loop shaping. **(b)** Crossover frequencies between selected noise contributions as a function of effective ADC resolution (ENOB). The blue curve indicates the frequency at which ADC noise equals actuator/feedback noise, demonstrating the systematic extension of the near-thermal bandwidth with increasing ENOB. The orange curve shows the crossover between readout and actuator noise, which remains nearly invariant with ENOB, reflecting its weak dependence on digitization resolution. Together, panels **(a)** and **(b)** illustrate how digital resolution primarily controls the high-frequency noise budget, while low- and mid-band performance are constrained by analog sensing and actuation mechanisms.

A primary metric is the minimum self-noise level, defined as

$$N_{\mathrm{a,min}} = \min_{\omega} N_{\mathrm{a}}(\omega), \tag{9}$$

where $N_{\mathrm{a}}(\omega)$ is the self-noise ASD defined in Eq. (4). While $N_{\mathrm{a,min}}$ provides a compact measure of ultimate sensitivity, it does not capture the frequency extent over which this sensitivity is achieved. As such, it should not be used in isolation to rank sensor designs.

Complementary information is provided by the near-plateau bandwidth introduced in Sect. 6, which quantifies the usable frequency interval over which the self-noise remains close to its minimum value. Together, these two metrics define a sensitivity–bandwidth trade-off that is central to sensor design.

To assess how closely a given design approaches the fundamental thermal noise limit, the margin to the thermal bound is evaluated using the ratio defined in Eq. (6). For practical design optimisation, a scalar metric can be introduced by averaging this ratio over a frequency interval of interest $[\omega_1, \omega_2]$,

$$\overline{R} = \frac{1}{\omega_2 - \omega_1} \int_{\omega_1}^{\omega_2} \frac{S_{\mathrm{a,self}}(\omega)}{S_{\mathrm{a,th}}(\omega)} \mathrm{d}\omega. \tag{10}$$

Values of $\overline{R}$ close to unity indicate near-thermal-limited performance over the selected band, whereas larger values highlight the dominance of implementation-dependent noise sources. This metric is particularly useful for identifying whether further optimisation efforts should focus on mechanical design or on technological improvements in readout and digital acquisition.

Beyond absolute performance, robustness with respect to parameter uncertainty constitutes a critical aspect of sensor design. Within the digital-twin framework, robustness is assessed by analysing the variability of the performance metrics introduced above under the assumed parameter uncertainty distributions.

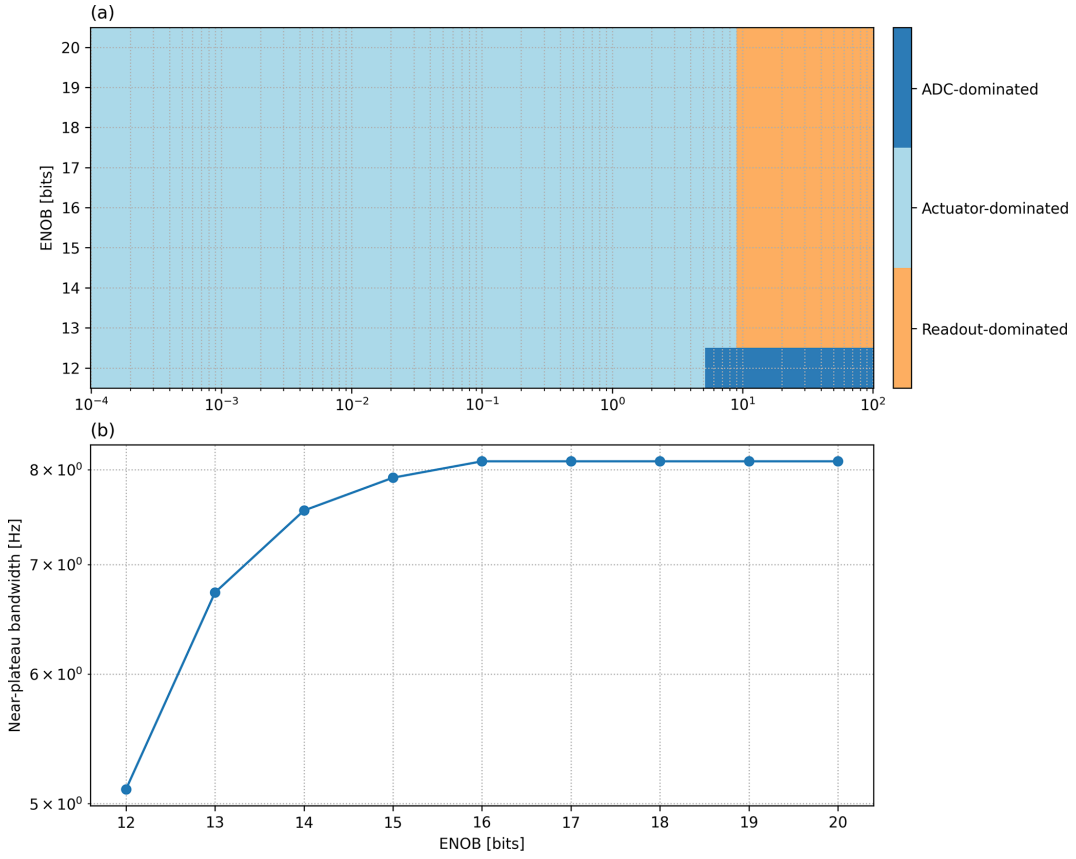

Figure 7. Design regime map and effective near-plateau bandwidth. **(a)** Dominant self-noise contributor in the ENOB–frequency plane, classified by variance comparison between readout, actuator/feedback, and ADC noise terms. The map highlights distinct operating regimes and shows that, for typical design parameters, the system is predominantly actuator-limited over most of the seismic bandwidth, with ADC-dominated behavior confined to low ENOB and high frequencies, and readout domination emerging at the highest frequencies. **(b)** Near-plateau bandwidth as a function of ENOB, defined as the frequency range over which the total self-noise remains within 30 % of its minimum achievable plateau. Increasing ENOB significantly extends the usable low-noise bandwidth up to ENOB $\approx$ 15–16, beyond which further improvements yield diminishing returns, indicating a transition to actuator/readout-limited performance. This representation provides a compact and quantitative guideline for digitization requirements in force-feedback broadband sensors.

Designs that achieve marginally lower self-noise but exhibit strong sensitivity to parameter variations may be less desirable than slightly noisier but more robust configurations. This consideration is especially relevant for long-term geoscientific deployments, where environmental variability and ageing effects can significantly impact performance (Wright and Davidson, 2024).

The optimisation strategy adopted here therefore prioritises uncertainty-aware performance, favouring regions of the design space in which sensitivity, bandwidth, and robustness are jointly optimised. The regime maps shown in Fig. 7 illustrate how such regions can be identified and compared across alternative design choices.

Within the proposed framework, design optimisation proceeds iteratively. Initial design choices define a candidate digital twin, whose performance metrics are evaluated using the methods described above. Identified bottlenecks then guide targeted modifications to mechanical parameters, control architecture, or digital acquisition settings, and the analysis is repeated until the desired performance objectives are met.

This optimisation process is methodologically consolidated, as it relies on physically interpretable metrics derived directly from the causal digital twin. At the same time, specific optimisation outcomes remain context-dependent, reflecting application-specific requirements and constraints. The broader implications of this balance between generality and specificity are discussed in Sect. 8.

8 Discussion

Recent advances in digital-twin methodologies for metrology have emphasised the importance of uncertainty quantification, traceability, and physically consistent virtual experimentation (Wright and Davidson, 2024; Maculotti et al.,

2024). The present framework aligns with these developments by extending such principles to force-feedback inertial sensors and explicitly incorporating causality constraints within the modeling architecture.

This section discusses the implications of the proposed causal and uncertainty-aware digital-twin framework, placing the results presented in Sects. 4–7 in a broader methodological and instrumental context. The discussion focuses on the generality of the approach, its limitations, and its relevance for the design of ultra–low-noise geoscientific inertial sensors.

A conceptual synthesis of the workflow and its role in guiding design decisions is illustrated in Fig. 8, which integrates the individual analysis steps into a unified methodological framework.

A key outcome of this study is that enforcing physical causality and realisability constraints at the digital-twin level is essential for obtaining meaningful metrological predictions. Non-causal or idealised representations may lead to overly optimistic noise estimates and obscure the distinction between fundamental and implementation-dependent limits. (Wright and Davidson, 2024; Maculotti et al., 2024).

The practical relevance of these metrics can be illustrated through representative design scenarios based on Figs. 4–7.

In a first scenario, consider a design operating close to the thermal noise limit over a restricted frequency band, as indicated by $R(\omega) \approx 1$ in Fig. 4. In this regime, further improvements in sensitivity cannot be achieved through readout or digital optimisation alone, and require modifications of the mechanical plant, such as reducing dissipation or increasing effective mass. This identifies the thermal-to-readout crossover as a key decision threshold between mechanical and electronic optimisation strategies.

In a second scenario, Fig. 6 shows that increasing ADC resolution shifts the crossover between digital and actuation noise toward higher frequencies, effectively extending the near-plateau bandwidth. However, as illustrated in Fig. 7, this improvement saturates beyond a certain ENOB (Effective Number of Bits), beyond which actuator or readout noise becomes dominant. This demonstrates that digital optimisation alone cannot ensure broadband performance gains.

A third scenario concerns the role of damping, as shown in Fig. 5. Variations in damping primarily affect low-frequency performance through their impact on thermal noise and feedback dynamics, while high-frequency behaviour remains largely unchanged. This indicates that mechanical and digital parameters act on distinct spectral regions, enabling targeted and decoupled optimisation strategies.

The performance metrics introduced in Sects. 6 and 7 provide a compact representation of complex spectral information and facilitate comparison between alternative sensor architectures. In particular, crossover frequencies and near-plateau bandwidths offer actionable indicators of where design effort is most effectively directed.

The proposed framework can be further contextualised by comparison with conventional sensor design methodologies. Traditional noise budgeting approaches typically evaluate individual noise contributions independently and combine them a posteriori, often neglecting system-level coupling effects introduced by feedback control and digital acquisition. While such approaches are effective for first-order performance estimation, they may lead to inconsistencies when applied to closed-loop architectures, where transfer functions are inherently interdependent.

Simplified digital twin models, on the other hand, often represent individual subsystems or rely on non-causal inversions to estimate idealised performance limits. Although computationally efficient, these formulations may overlook constraints imposed by causality and realisability, leading to optimistic predictions of achievable sensitivity.

In contrast, the present framework enforces causal system dynamics at all stages and propagates all noise sources through a unified set of transfer functions, enabling a consistent system-level evaluation. The resulting performance metrics, including crossover frequencies (Eq. 7) and near-plateau bandwidth (Eq. 8), provide quantitative indicators that are not directly accessible through conventional noise budgeting approaches.

From a quantitative perspective, the ratio to the thermal limit (Eq. 6) and its band-averaged form (Eq. 10) enable a direct comparison between designs, while explicitly accounting for implementation-dependent constraints. These metrics highlight how different design strategies shift regime boundaries rather than uniformly improving performance, an effect that is not captured by traditional component-wise optimisation.

This comparison indicates that the main advantage of the proposed digital twin lies not in redefining fundamental limits, but in providing a physically consistent and uncertainty-aware framework to approach them in a controlled and interpretable manner.

A schematic comparison of the different methodological approaches is provided in Fig. 9, highlighting how the proposed framework enables physically consistent and uncertainty-aware system-level design.

The regime maps derived from these metrics should be interpreted as decision-support tools rather than as absolute performance predictors. While the existence of distinct noise-dominated regimes is a robust outcome of the noise-budget formalism, the precise boundaries between regimes depend on modelling assumptions and uncertainty characterisation. This sensitivity underscores the importance of uncertainty-aware optimisation strategies in sensor design.

Although the digital twin is formulated with force-feedback inertial sensors in mind, the underlying methodology is not restricted to a specific sensor type. The same framework can be adapted to other classes of geoscientific instruments, such as tiltmeters, gravimeters, or strain sensors, including atom-interferometry gravimeters demonstrated in

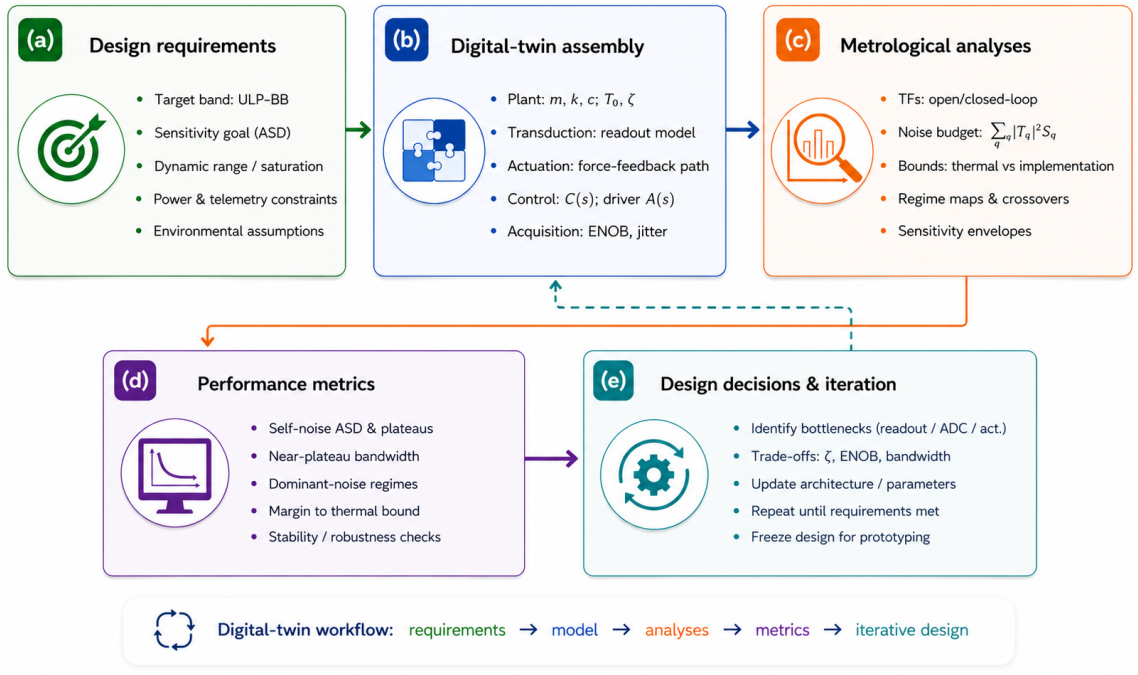

Figure 8. Conceptual workflow of the digital-twin framework adopted for the engineering design and metrological assessment of ultra-low-noise broadband inertial sensors. **(a)** Definition of design requirements, including target bandwidth, sensitivity objectives, dynamic range constraints, and environmental assumptions. **(b)** Assembly of the digital twin, integrating the mechanical plant, transduction chain, force-feedback actuation, control architecture, and acquisition system. **(c)** Metrological analyses performed in the frequency domain, encompassing open- and closed-loop transfer functions, full noise-budget decomposition, and theoretical versus implementation-dependent performance bounds. **(d)** Extraction of performance metrics, such as self-noise plateaus, near-plateau bandwidth, dominant noise regimes, and margins relative to the thermal limit. **(e)** Iterative design loop, where identified bottlenecks guide trade-offs among damping, ENOB, and bandwidth, leading to progressive architectural refinement and final design freeze. Solid arrows indicate the nominal forward workflow, whereas the dashed arrow highlights the feedback loop driving iterative optimization.

harsh volcanic environments (Antoni-Micollier et al., 2022), provided that their dynamics can be represented within a linear, causal system framework.

The abstraction of transduction, actuation, and digital acquisition as modular subsystems facilitates such extensions and supports the development of sensor-specific digital twins within a common metrological structure. This generality is consistent with recent efforts to standardise digital-twin methodologies across measurement domains (Ríos et al., 2020; Vitale et al., 2022; Ubezio et al., 2023).

Several limitations of the present framework should be acknowledged. First, the analysis assumes linearity, stationarity, and mutual incoherence of noise sources. While these assumptions are standard in inertial sensor metrology, they may be violated under extreme operating conditions or in the presence of strong environmental coupling.

Second, parameter uncertainties are treated as statistically independent unless explicitly stated. In real instruments, correlations between mechanical, electronic, and environmental parameters may exist and could affect uncertainty propagation. Incorporating such correlations would require additional modelling effort and data, which lies beyond the scope of the present study.

Finally, the framework does not incorporate site-specific noise conditions or experimental calibration data (D'Alessandro et al., 2021; Catania et al., 2025). As a result, the predicted performance should be interpreted as an intrinsic sensor capability rather than as a guarantee of field performance. This distinction is particularly important when comparing digital-twin predictions with observational data.

A further limitation of the present study is the lack of direct experimental validation against specific sensor prototypes. This reflects the intended role of the proposed digital twin as a design-stage and feasibility-assessment tool, rather than as a post-deployment calibration framework.

In this context, the objective is not to reproduce the performance of a particular instrument, but to provide a physically consistent and causality-constrained environment for exploring design trade-offs and identifying fundamental and implementation-dependent performance limits.

Nevertheless, the integration of the digital twin with experimental data represents a natural and important extension of the present work, as demonstrated by recent metrology-

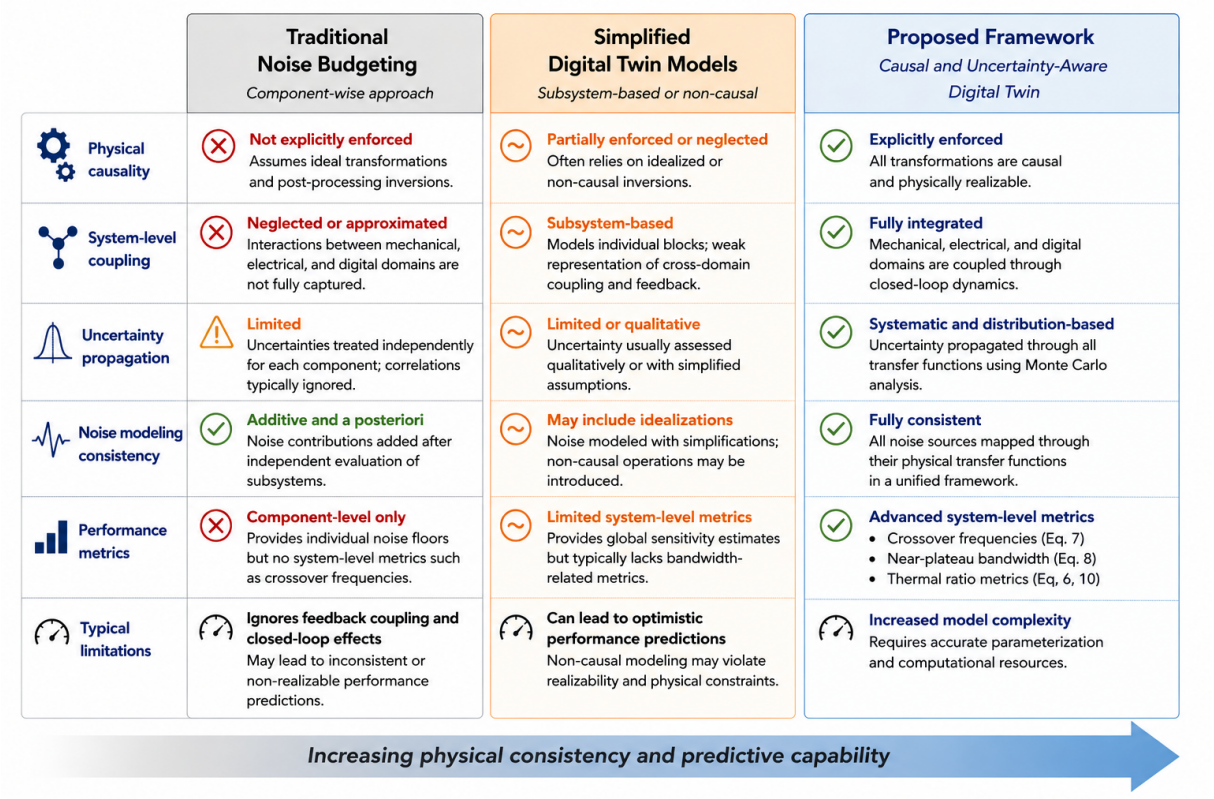

Figure 9. Comparative assessment of sensor design methodologies for ultra–low-noise inertial systems. The diagram contrasts traditional noise budgeting, simplified digital twin approaches, and the proposed causal and uncertainty-aware digital-twin framework in terms of physical causality, system-level coupling, uncertainty propagation, noise modeling consistency, and performance metrics. Conventional approaches either neglect feedback-induced coupling or rely on simplified or non-causal representations, limiting their ability to provide consistent system-level predictions. In contrast, the proposed framework enforces causal dynamics and propagates all noise sources through a unified transfer-function representation, enabling physically consistent evaluation and the definition of advanced performance metrics such as crossover frequencies and near-plateau bandwidth.

integrated digital twin frameworks that explicitly combine physical models with measurement data for validation and calibration purposes (Samadi et al., 2025). Future developments may include the use of laboratory or field measurements to calibrate model parameters, validate predicted self-noise spectra, and refine uncertainty distributions within a data-informed framework.

Despite these limitations, the proposed framework provides a solid foundation for future extensions. Potential developments include the incorporation of nonlinear effects, the integration of experimental calibration data to refine parameter distributions, and the coupling of the digital twin with real-time monitoring systems. Data-driven denoising approaches have also been explored for inertial sensors/accelerometers and may complement model-based digital twins (Yang et al., 2023).

Such extensions would further enhance the utility of digital twins as tools for both design-stage optimisation and operational performance assessment in geoscientific instrumentation.

9 Conclusions

This work introduces a causal and uncertainty-aware digital-twin framework for the design and metrological assessment of ultra–low-noise geoscientific inertial sensors. The framework integrates mechanical dynamics, force-feedback control, transduction, and digital acquisition within a physically realisable and causally consistent representation.

A central contribution is the formulation of a unified noise-budget model that distinguishes between fundamental thermal limits and implementation-dependent noise sources, enabling physically interpretable self-noise predictions and uncertainty-aware performance assessment.

The results demonstrate that enforcing causality and realisability is essential for obtaining reliable performance estimates, particularly in ultra–low-noise regimes where idealised formulations may lead to overly optimistic predictions.

Despite the simplifying assumptions adopted, the proposed framework provides a robust and extensible foundation for design-stage optimisation and virtual experimenta-

tion. Potential extensions include the incorporation of nonlinear effects, correlated uncertainties, and experimental calibration data.

Beyond inertial sensors, the methodology is readily transferable to other classes of geoscientific instruments, supporting the broader adoption of digital twins as design and decision-support tools in geoscientific instrumentation.

Code availability. The custom Python scripts used to reproduce the computational figure-generation workflow associated with this study are publicly archived on Zenodo at https://doi.org/10.5281/zenodo.21155578 (D'Alessandro, 2026). The archived repository contains the standalone scripts used to generate Figs. 2–7, together with a README file, software requirements, licensing information, and citation metadata. Figures 1, 8, and 9 are conceptual graphical schematics prepared separately and are not generated by the computational scripts. All custom software required to reproduce the numerical analyses presented in the manuscript is contained in a single repository. No third-party source code is redistributed in the archive; the scripts rely only on standard open-source scientific Python libraries documented in the repository metadata.

Data availability. No data sets were used in this article.

Competing interests. The author has declared that there are no competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. The authors bear the ultimate responsibility for providing appropriate place names. Views expressed in the text are those of the authors and do not necessarily reflect the views of the publisher.

Review statement. This paper was edited by Alessandro Fedeli and reviewed by two anonymous referees.

References

Antoni-Micollier, L., Carbone, D., Ménoret, V., Lautier-Gaud, J., King, T., and Greco, F.: Detecting volcano-related underground mass changes with a quantum gravimeter, Geophys. Res. Lett., 49, e2022GL097814, https://doi.org/10.1029/2022GL097814, 2022.

Bai, W., Feng, W., Wang, P., Zhang, Z., and Zhao, G.: Research on Interferometric Tilt Sensor for Vibration Isolation Platform, Sensors-Basel, 25, 1777, https://doi.org/10.3390/s25061777, 2025.

Bendat, J. S. and Piersol, A. G.: Random Data: Analysis and Measurement Procedures, 4th edn., Wiley, https://doi.org/10.1002/9781118032428, 2010.

Brown, D. D., Jones, P., Rowlinson, S., Leavey, S., Green, A. C., Töyrä, D., and Freise, A.: Pykat: Python package for modelling precision optical interferometers, SoftwareX, 12, 100613, https://doi.org/10.1016/j.softx.2020.100613, 2020.

Carlin, H. M., Goodall, P. A., Young, R. I. M., and West, A. A.: An interactive framework to support decision-making for Digital Twin design, Journal of Industrial Information Integration, 41, 100639, https://doi.org/10.1016/j.jii.2024.100639, 2024.

Carter, J. J., Birckigt, P., Gerberding, O., and Koehlenbeck, S. M.: High precision inertial sensors on a one inch diameter optic, Sci. Rep.-UK, 14, https://doi.org/10.1038/s41598-024-68623-0, 2024.

Carter, J. J., Birckigt, P., Lehmann, J., Basalaev, A., Kranzhoff, L., Al-Kershi, S., Carlassara, M., Chiarini, G., Khan, F., and Leibeling, G.: Testing compact, fused silica resonator based inertial sensors in a gravitational wave detector prototype facility, Classical Quant. Grav., 42, 185001, https://doi.org/10.1088/1361-6382/adff34, 2025.

Catania, M., Figlioli, A., Vitale, G., and D'Alessandro, A.: Seismic noise characterization of broad-band stations in the Italian region using power spectral density: a frequency, spatial and statistical analysis, Geophys. J. Int., 242, ggaf168, https://doi.org/10.1093/gji/ggaf168, 2025.

Chen, Y., Liu, X., Wang, L., Yu, T., Wang, Z., Xue, K., Sui, Y., and Chen, Y.: Research and optimization of high-performance front-end circuit noise for inertial sensors, Sensors-Basel, 24, 805, https://doi.org/10.3390/s24030805, 2024.

Collette, C., Janssens, S., Fernandez-Carmona, P., Artoos, K., Guinchard, M., Hauviller, C., and Preumont, A.: Review: Inertial sensors for low-frequency seismic vibration measurement, B. Seismol. Soc. Am., 102, 1289–1300, https://doi.org/10.1785/0120110223, 2012.

D'Alessandro, A.: Software code for "Causal and uncertainty-aware digital-twin framework for ultra-low-noise geoscientific inertial sensors", Zenodo [code], https://doi.org/10.5281/zenodo.21155578, 2026.

D'Alessandro, A., Scudero, S., and Vitale, G.: A Review of the Capacitive MEMS for Seismology, Sensors-Basel, 19, 3093, https://doi.org/10.3390/s19143093, 2019.

D'Alessandro, A., Greco, L., Scudero, S., and Lauciani, V.: Spectral Characterization and Spatiotemporal Variability of the Background Seismic Noise in Italy, Earth and Space Science, 8, e2020EA001579, https://doi.org/10.1029/2020EA001579, 2021.

Ding, B., Zhao, G., Watchi, J., Sider, A., and Collette, C.: An interferometric inertial sensor for low-frequency seismic isolation, Sensor. Actuat. A-Phys., 335, 113398, https://doi.org/10.1016/j.sna.2022.113398, 2022.

El-Sheimy, N., Yoon, S., and Jiang, Z.: Inertial sensors technologies and error modelling for navigation and uncertainty analysis, Journal of Applied Geodesy, 14, 115–134, https://doi.org/10.1186/s43020-019-0001-5, 2020.

Kranzhoff, S. L., Lehmann, J., Kirchhoff, R., Carlssara, M., Cooper, S. J., Koch, P., Leavey, S., Lück, H., Mow-Lowry, C. M., and Wöhler, J.: A vertical inertial sensor with interferometric readout, Classical Quant. Grav., 40, 015007, https://doi.org/10.1088/1361-6382/aca45b, 2023.

Maculotti, G., Marschall, M., Kok, G., Chekh, B. A., van Dijk, M., Flores, J., Genta, G., Puerto, P., Galetto, M., and Schmelter, S.: A Shared Metrological Framework for Trustworthy Virtual Experiments and Digital Twins, Metrology, 4, 337–363, https://doi.org/10.3390/metrology4030021, 2024.

Mitchell, A., Lehmann, J., Koch, P., Cooper, S. J., van Dongen, J., Prokhorov, L., Holland, N. A., Valentini, M., Saffarieh, P., and Mow-Lowry, C. M.: Integration of high-performance compact interferometric sensors in a suspended interferometer, Classical Quant. Grav., 42, 195014, https://doi.org/10.1088/1361-6382/ae0087, 2025.

Oppenheim, A. V. and Schafer, R. W.: Discrete-Time Signal Processing, 3rd edn., Prentice Hall, ISBN: 9780131988424, 2010.

Poroskun, I., Rothleitner, C., and Heißelmann, D.: Structure of digital metrological twins as software for uncertainty estimation, J. Sens. Sens. Syst., 11, 75–82, https://doi.org/10.5194/jsss-11-75-2022, 2022.

Prasad, A., Middlemiss, R. P., Noack, A., Anastasiou, K., Bramsiepe, S. G., Toland, K., Utting, P. R., Paul, D. J., and Hammond, G. D.: A 19 d earth tide measurement with a MEMS gravimeter, Sci. Rep.-UK, 12, 13091, https://doi.org/10.1038/s41598-022-16881-1, 2022.

Ríos, J., Staudter, G., Weber, M., and Anderl, R.: Uncertainty of data and the digital twin: a review, International Journal of Product Lifecycle Management, 12, 329–358, https://doi.org/10.1504/IJPLM.2020.10035102, 2020.

Samadi, H., Ahsan, M. M., and Raman, S.: Metrology and manufacturing-integrated digital twin (MM-DT) for advanced manufacturing: Insights from coordinate measuring machine (CMM) and FARO arm measurements, Next Research, 2, 100299, https://doi.org/10.1016/j.nexres.2025.100299, 2025.

Scholz, G., Fortmeier, I., Marschall, M., Stavridis, M., Schulz, M., and Elster, C.: Experimental design for virtual experiments in tilted-wave interferometry, Metrology, 2, 84–97, https://doi.org/10.3390/metrology2010006, 2022.

Sepahi-Boroujeni, S. and Khameneifar, F.: Digital twin-enabled error and uncertainty mapping for 3D scanning, Precis. Eng., https://doi.org/10.1016/j.precisioneng.2024.03.007, 2024.

Sider, A., Di Fronzo, C., Amez-Droz, L., Amorosi, A., Badaracco, F., Baer, P., Bertolini, A., Bruno, G., Cebeci, P., Collette, C., Ebert, J., Erben, B., Esteves, R., Ferreira, E., Gatti, A., Giesberts, M., Hebbeker, T., van Heijningen, J., Hennig, J.-S., Hennig, M., Hild, S., Hoefer, M., Hoffmann, H.-D., Jacques, L., Jamshidi, R., Joppe, R., Kuhlbusch, T.-J., Lakkis, M. H., Lenaerts, C., Locquet, J.-P., Loicq, J., Le Van, B. L., Loosen, P., Nesladek, M., Reiter, M., Stahl, A., Steinlechner, J., Steinlechner, S., Tavernier, F., Teloi, M., Villaboa Perez, J., and Zeoli, M.: E-TEST: a compact low-frequency isolator for a large cryogenic mirror, Classical Quant. Grav., 40, 165002, https://doi.org/10.1088/1361-6382/ace230, 2023.

Sun, Z., Jayasinghe, S., Sidiq, A., Shahrivar, F., Mahmoodian, M., and Setunge, S.: Approach towards the development of digital twin for structural health monitoring of civil infrastructure: a comprehensive review, Sensors-Basel, 25, 59, https://doi.org/10.3390/s25010059, 2025.

Ubezio, B., Ergun, S., and Zangl, H.: Realistic sensor simulations for the digital twin, e+i Elektrotechnik und Informationstechnik, 140, 562–571, https://doi.org/10.1007/s00502-023-01156-y, 2023.

Ubhi, A. S., Prokhorov, L., Cooper, S., Di Fronzo, C., Bryant, J., Hoyland, D., Mitchell, A., van Dongen, J., Mow-Lowry, C., Cumming, A., Hammond, G., and Martynov, D.: Active platform stabilization with a 6D seismometer, Appl. Phys. Lett., 121, 174101, https://doi.org/10.1063/5.0118606, 2022.

van Dongen, J., Prokhorov, L., Cooper, S. J., Barton, M. A., Bonilla, E., Dooley, K. L., Driggers, J. C., Effler, A., Holland, N. A., Huddart, A., Kasprzack, M., Kissel, J. S., Lantz, B., Mitchell, A. L., O'Dell, J., Pele, A., Robertson, C., and Mow-Lowry, C. M.: Reducing control noise in gravitational wave detectors with interferometric local damping of suspended optics, Rev. Sci. Instrum., 94, 054501, https://doi.org/10.1063/5.0144865, 2023.

Vitale, G., D'Alessandro, A., Di Benedetto, A., Figlioli, A., Costanzo, A., Speciale, S., Piattoni, Q., and Cipriani, L.: Urban Seismic Network Based on MEMS Sensors: The Experience of the Seismic Observatory in Camerino (Marche, Italy), Sensors-Basel, 22, 4335, https://doi.org/10.3390/s22124335, 2022.

Vlaeyen, M., Haitjema, H., and Dewulf, W.: Digital twin of an optical measurement system, Sensors-Basel, 21, 6638, https://doi.org/10.3390/s21196638, 2021.

Wright, L. and Davidson, S.: Digital twins for metrology; metrology for digital twins, Meas. Sci. Technol., 35, https://doi.org/10.1088/1361-6501/ad2050, 2024.

Wübbeler, G., Marschall, M., Kniel, K., Heißelmann, D., Härtig, F., and Elster, C.: GUM-Compliant Uncertainty Evaluation Using Virtual Experiments, Metrology, 2, 114–127, https://doi.org/10.3390/metrology2010008, 2022.

Yang, Z., Zhang, H., Xu, P., and Luo, Z.: Unsupervised Noise Reductions for Gravitational Reference Sensors or Accelerometers Based on the Noise2Noise Method, Sensors-Basel, 23, 6030, https://doi.org/10.3390/s23136030, 2023.