AIR Journal of Engineering and Technology, Vol. 2026
Article Unique ID: AIRJET2026811
DOI: 10.65737/AIRJET2026811

Why 7G Needs New Theorems: A Geodesy-Based Perspective on Beyond-6G Foundations

Mosab Hawarey
Director, Geospatial Research (*director@geospatial.ch*)

Received: 2026 June 07 | Revised: 2026 June 15 | Approved: 2026 June 16 | Published: 2026 June 19

Abstract

Every prior generation of cellular wireless was preceded by a foundational theorem: Shannon for the digital transition of 1G–3G, Telatar for the multi-antenna capacity that defined 4G, Arıkan for the polar codes that underpin 5G control signaling. Sixth-generation wireless is deep into its standardization cycle without an analogous defining result, and seventh-generation wireless has not yet been posed as a theoretical problem at all. We argue, from the perspective of a geodetic engineer who has spent two decades working with the original integrated sensing and communication system—the Global Navigation Satellite System—that the most consequential contributions to 7G will not come from faster hardware but from new theorems, and that the geodetic estimation tradition provides the most mature and most immediately applicable toolkit for producing them. We identify five concrete mathematical bridges from geodesy to beyond-6G open problems: Cramér–Rao analysis to ISAC capacity–distortion regions, ambiguity-function design to delay-Doppler channel theory, multi-constellation estimation to Space-Air-Ground Integrated Network capacity, Kalman-filter convergence to AI-native channel tracking with provable guarantees, and antenna phase-center calibration to near-field electromagnetic information theory. We frame the 2026–2030 window—bounded by the ITU-R IMT-2030 evaluation phase, World Radiocommunication Conference 2027 spectrum decisions, and the first 6G releases of 3GPP—as the critical period in which foundational theoretical work can still shape what systems get built, and issue a call to the geodesy and wireless-communications research communities to engage across the boundary that currently separates them.

Keywords: 7G; beyond-6G; integrated sensing and communication (ISAC); Cramér–Rao bound; Global Navigation Satellite System (GNSS); Kalman filtering; delay-Doppler channel; Space-Air-Ground Integrated Network (SAGIN); AI-native physical layer; electromagnetic information theory.

Citation: Hawarey, M. (2026). Why 7G Needs New Theorems: A Geodesy-Based Perspective on Beyond-6G Foundations. *AIR Journal of Engineering and Technology*, Vol. 2026, AIRJET2026811. https://doi.org/10.65737/AIRJET2026811

1. Introduction

What is 6G's theorem? Shannon's 1948 coding theorem [1] supplied the objective function for the analog-to-digital transition of first through third generation systems; Telatar's characterization of multi-antenna capacity [2] converted MIMO from a clever engineering idea to the defining technical feature of 4G; Arıkan's channel polarization [3] gave 5G its first explicit capacity-achieving code family. Each generation was preceded by a single citable result that constrained what the generation could become. Sixth-generation wireless, now deep into its standardization cycle, has no such result. Seventh-generation wireless has not yet been posed as a theoretical problem at all.

Sixth generation wireless is, by most visible measures, on schedule. The ITU-R IMT-2030 framework is published [4]; World Radiocommunication Conference 2027 will issue candidate spectrum allocations; 3GPP Release 21 is slated to contain the first normative 6G specification. What is missing is the theorem. There is no Shannon for 6G, no Telatar, no Arıkan. The ISAC capacity–distortion region, the delay-Doppler channel capacity, the provable-guarantee analog of the Kalman covariance for AI-native physical layers—each is the subject of a rapidly expanding literature [7]–[10] for ISAC and [11] for delay-Doppler waveforms, and none has been established in the generality needed to constrain the generation. If 6G lacks a defining theorem, seventh generation wireless lacks even the problem statement: current uses of the term "7G" in the literature are speculative or use-case driven [12], [13], never theoretical.

This paper argues that beyond-6G wireless requires a distinct layer of foundational mathematical work that the community has not yet begun, and that the geodetic-engineering tradition—specifically Cramér–Rao analysis, Kalman estimation theory, ambiguity-function design, and multi-constellation integration—provides the missing toolkit. The argument rests on a recognition rather than a claim: the Global Navigation Satellite System [14], [15] has, for fifty years, been a deployed integrated sensing and communication system at planetary scale, and the mathematical apparatus that makes it work [16]–[22] addresses problems the 6G ISAC community currently treats as open. We identify five concrete mathematical bridges from geodesy to beyond-6G open problems—from Cramér–Rao to ISAC capacity–distortion regions, from ambiguity functions to delay-Doppler channel theory, from multi-constellation estimation to SAGIN capacity, from Kalman filtering to AI-native channel tracking with provable guarantees, and from antenna phase-center calibration to near-field electromagnetic information theory. We frame the 2026–2030 window as the critical period in which foundational theoretical work can still shape what systems get built, and issue a call to the geodesy and wireless-communications research communities to engage across the boundary that currently separates them.

The remainder of the paper is organized as follows. Section 2 establishes the historical pattern of theorems preceding wireless generations and the anomaly that 6G and 7G present to that pattern. Section 3 argues that GNSS is the original ISAC system and that its mathematical apparatus is structurally identical to what beyond-6G theory is currently seeking. Section 4 presents the five mathematical bridges. Section 5 makes the timing case anchored in ITU, WRC, and 3GPP calendars. Section 6 issues the call to action.

2. The Pattern: Theorems Before Systems

Every mature technology rests on mathematics that was in place before the technology became inevitable. Wireless communication is no exception. The digital-communications era begins in 1948 with Claude Shannon's demonstration that a noisy channel has a well-defined capacity and that this capacity can be approached by coding [1], [49]—a result without which the analog-to-digital transition that enabled first through third generation cellular systems would have lacked an objective function. Forty years later, Telatar's derivation of the capacity of multiple-antenna channels [2] gave fourth generation wireless its defining technical feature: MIMO transmission moved from a clever engineering idea to a provably optimal use of spatial degrees of freedom. Fifth generation systems, in turn, were preceded by Arıkan's introduction of polar codes [3], which provided the first explicit family of codes proven to achieve capacity on binary-input symmetric channels and now serve as the control-channel coding in the 3GPP 5G New Radio standard. This is a pattern worth naming. In each case the mathematics preceded the standards work by a decade or more; in each case the theorem was narrow enough to be proven and general enough to constrain an entire generation of hardware; and in each case practitioners could point to a single, citable result and say: this is why the system takes the form it does. The pattern is not universal—many engineering advances have proceeded without a governing theorem—but for the cellular generations the regularity has held. A reasonable question, then, is whether it will hold again.

Two caveats temper the pattern before it is extended. First, the inductive base is three transitions, not a large sample; it is, however, the entire population of completed generational transitions to date rather than a subset of a larger one, so the regularity is at least exhaustive of the available cases. Second, the pattern presumes that a generation is organized around a single governing result. That presumption may not hold beyond 5G: sixth- and seventh-generation wireless may prove inherently multi-theoretic, defined by a small family of results—one for the sensing–communication tradeoff, one for time-varying-channel capacity, one for learning-based estimation with guarantees—rather than by one Shannon-style theorem. The argument of this paper does not depend on the single-theorem reading. The five bridges of Section 4 are individually self-contained, and the call they motivate stands whether the foundations of 7G turn out to be one theorem or several.

Against this backdrop, the current state of sixth generation research is striking. The International Telecommunication Union's IMT-2030 framework [4] describes six usage scenarios and sets aspirational targets for throughput, latency, coverage, and positioning accuracy; the 6G literature enumerates enabling technologies—terahertz propagation, ultra-massive MIMO, reconfigurable intelligent surfaces, non-terrestrial networks, AI-native air interfaces [5], [6]—and industry roadmaps anchor commercial deployment around 2030, with World Radiocommunication Conference 2027 expected to issue candidate spectrum allocations and 3GPP Release 21 slated to contain the first normative 6G specifications. By any visible measure, sixth generation wireless is on schedule. Yet one item on the checklist is blank. There is no Shannon for 6G. There is no Telatar, no Arıkan. There is no single, citable theorem that constrains what a 6G system must look like the way capacity, MIMO capacity, and channel polarization constrained their respective generations. Integrated sensing and communication (ISAC), the most theoretically ambitious of the 6G pillars, is the subject of a rapidly expanding literature on waveform design, beamforming, and network-level throughput–sensing tradeoffs [7]–[10], but a unified capacity–distortion region—the analog of Shannon's curve—remains unproven in the general case. Delay-Doppler channel theory, now central to the orthogonal time-frequency-space modulation family [11], has produced no counterpart to Shannon's channel coding theorem. AI-native physical layers are being deployed faster than convergence or generalization guarantees are being derived for them. The machinery required to answer the fundamental questions exists in pieces; it has not been assembled.

If 6G lacks a defining theorem, seventh generation wireless lacks even the problem statement. The term "7G" currently circulates in two contexts: speculative end-of-survey sections appended to 6G visions, and use-case-oriented pieces describing entrepreneurial opportunities or multi-path transport protocols [12], [13]. In neither context has 7G been posed as a distinct theoretical object. No ITU study group, no 3GPP work item, no flagship research program treats 7G as a question whose answer requires new mathematics. This is not a claim that 7G does not need new mathematics; it is a claim that the community has not yet asked. The central argument of this paper is that the community should ask, that the 2026–2030 window is the appropriate time to begin asking, and that geodetic engineering—the mathematical tradition that has underwritten precision positioning for half a century—offers a larger share of the answer than is currently recognized.

3. GNSS: The Original ISAC System

A system in which a single radio signal simultaneously carries data and permits the receiver to estimate range, velocity, and atmospheric state is normally described in the wireless literature as a candidate architecture for sixth or seventh generation networks. The description is apt except for the tense: such a system has existed for fifty years and is running on approximately a dozen active spacecraft over every head that reads this paper. The Global Positioning System, now one of four operational constellations within the broader Global Navigation Satellite System (GNSS) [14], [15], [47], transmits a navigation message (data), a ranging code (timing), and two carrier frequencies whose differential delay encodes ionospheric electron content (atmospheric sensing). A GNSS receiver performs, simultaneously and from the same electromagnetic observations, decoding, ranging, and environmental inference. This is the object that integrated sensing and communication theory is attempting to describe, and it predates the acronym. The mathematical apparatus that makes GNSS work is equally mature: the Cramér–Rao bound is the operational accuracy floor against which receiver designs have been benchmarked since the 1980s [16]–[19]; dilution of precision, a geometric rescaling of the CRB that separates signal-to-noise contributions from satellite-geometry contributions, is a standard output of every consumer receiver [14]; the extended Kalman filter, applied in tightly coupled GNSS–inertial integration, provides provably convergent state estimation and is the template on which every precision positioning product is built [20], [45]; carrier-phase ambiguity resolution via the LAMBDA method converts a mixed-integer estimation problem into a real-valued one with quantified success probability [21]; and atmospheric delay modeling [22] turns a propagation nuisance into a measurable environmental variable. The fact that detonations and earthquakes produce detectable travelling ionospheric disturbances in GNSS-derived total electron content [23], [24] is one of many demonstrations that this sensing modality is not a theoretical curiosity.

Two observations about this apparatus are relevant to wireless communications. The first is that the problems it solves are structurally identical to problems the ISAC community currently treats as open. The CRB that governs GNSS pseudorange estimation is the same Fisher-information object that must govern any ISAC system's sensing accuracy. The ambiguity function that GNSS signal designers have used to characterize code auto- and cross-correlation is the same two-dimensional delay-Doppler kernel that Zak-OTFS and related waveforms now place at the center of beyond-6G modulation theory [11]. The state-space formulation that supports Kalman-filtered positioning is directly applicable to channel tracking in mobile and non-terrestrial networks. None of this requires translation; it requires recognition. The second observation is that the geodetic-engineering community has a half-century head start on the hard instances of these problems—multi-constellation data fusion, centimeter-level precision under adversarial atmospheric conditions, integer-valued estimation at scale—and that this head start is not being drawn upon. The reason is largely sociological. Geodesy and wireless communications have separate conferences, separate journals, separate graduate curricula, and largely separate vocabularies for the same mathematics. The cost of this separation has been tolerable when the two fields' practical domains did not overlap. As 6G ISAC matures and as seventh generation wireless begins to be defined, the overlap becomes total. A field that has been doing integrated sensing and communication since the 1970s cannot remain outside the conversation about the theoretical foundations of integrated sensing and communication.

The present author writes from that field. Two decades of work by the author and collaborators on GNSS signal processing, ionospheric detection of explosive events using the Southern California GPS network [23], and observation of travelling ionospheric disturbances through GNSS total electron content analysis [24], supply the perspective offered here. The argument is not that geodesists should now do wireless research, nor that wireless researchers should now do geodesy. It is that certain specific mathematical bridges between the two fields are tight, load-bearing, and currently unused. The next section names five.

4. Five Bridges from Geodesy to 7G

The five bridges described below share a common structure. In each case a mathematical object familiar from geodesy—an accuracy bound, a correlation kernel, a multi-source fusion rule, a recursive estimator, a calibration model—has a twin in the beyond-6G wireless literature that is currently treated as an open problem or an ad hoc design choice. In each case the geodetic tradition has a half-century of deployed experience with the object and its failure modes, the tightness of the mapping can be stated in one paragraph, and the formal proof is the subject of a separate, forthcoming paper in this program's Phase 1 series. The bridges are presented in order of increasing speculativeness: the first is operational, the last remains conjectural, and the middle three sit in between. Figure 1 summarizes the mapping.

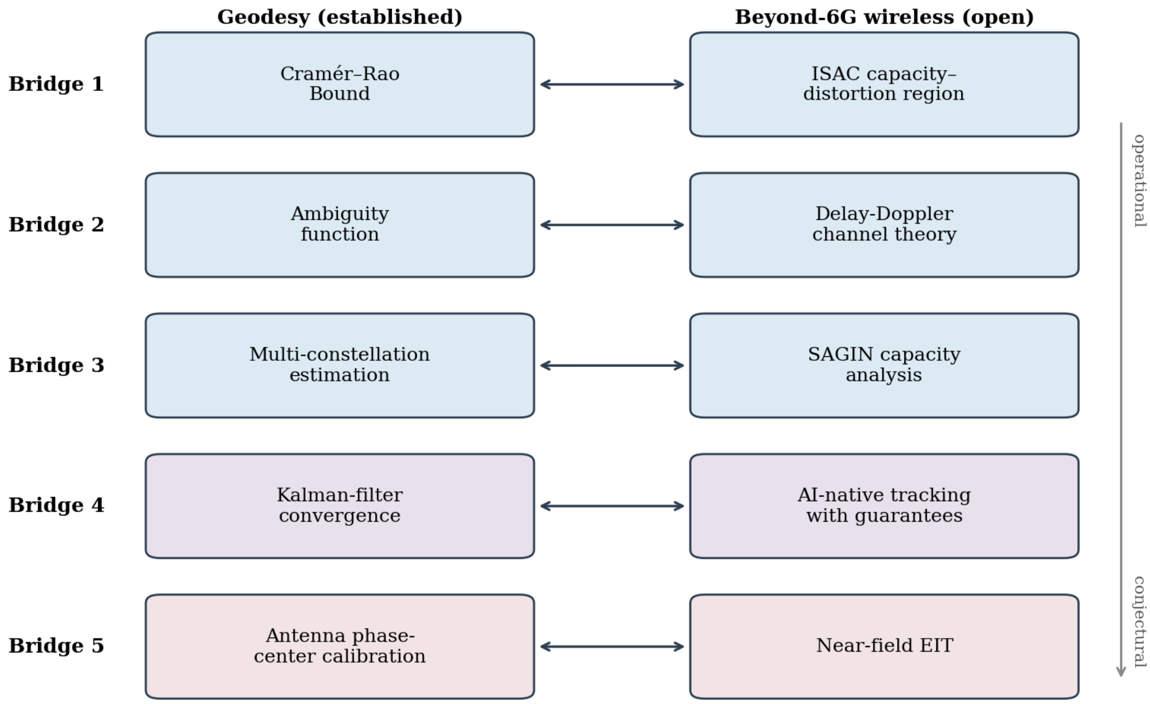

Figure 1. Five mathematical bridges from geodesy to beyond-6G wireless. Each row pairs an established geodetic concept (left) with a currently open wireless-communications problem (right). Bridges are ordered by increasing speculativeness: Bridge 1 is operational, supported by a first closed-form result [25]; Bridges 2–4 are concrete and staged for forthcoming theorem papers in the same 7G research program; Bridge 5 is a research direction rather than a proven link.

4.1 Bridge 1: The Cramér–Rao Bound and the ISAC Capacity–Distortion Region

The Cramér–Rao bound (CRB) states that the variance of any unbiased estimator of a parameter is lower-bounded by the inverse of the Fisher information available about that parameter in the observation [36]–[38]. For GNSS pseudorange estimation the bound is classical: the ranging variance attainable from a signal of root-mean-square bandwidth β and signal-to-noise ratio ρ is proportional to $(\beta^2\rho)^{-1}$, and the three-dimensional positioning covariance inherits this structure scaled by a geometry matrix whose singular values give the familiar dilution-of-precision quantities [14], [16]–[19]. Receiver designs—chip rate, correlator spacing, front-end filtering—are optimized against this bound, which is tight for high-SNR regimes with regular satellite geometry. An ISAC system performing simultaneous ranging and data transmission over a shared waveform faces the same estimation problem: the sensing-side performance floor is a Cramér–Rao bound on whatever environmental parameter the system is estimating—target range, target velocity, channel delay, angle of arrival—while the communication-side performance ceiling is a Shannon capacity. The two bounds are coupled because they share the same signal, the same power, the same time-frequency footprint, and often the same antenna aperture. The coupled object—the set of simultaneously achievable (rate, distortion) pairs—is the ISAC capacity–distortion region [7]–[10]. For restricted settings this region has been characterized; for the general multi-user, multi-target, fading-channel case it has not.

The geodetic contribution is twofold. First, the dilution-of-precision rescaling is not merely an engineering convenience; it is the right decomposition of the Fisher information matrix for any positioning-relevant ISAC design, and it cleanly separates the waveform-design problem (maximize β^2) from the geometry-design problem (minimize the condition number of the geometry matrix). Second, the high-SNR asymptotics that make GNSS receivers predictable—equivalence of maximum-likelihood and least-squares estimation, Gaussian error distributions, closed-form error propagation—apply unchanged to the ISAC sensing side in the regime where most deployed systems operate. A companion letter in this program [25] derives, in closed form, the CRB–Rate tradeoff for a GNSS-integrated ISAC link under AWGN, shows the Pareto frontier to be convex with an explicit optimal power allocation, and demonstrates formally that the classical GNSS geometric dilution of precision is a boundary case of this region. The present paper's claim is that this result is not a curiosity but the first operational instance of a bridge the community should now traverse in generality.

Box 1. Bridge 1 in brief: from the ISAC Fisher information to the GNSS CRB and its coupling to the Shannon rate.

Setup. A single shared waveform $s(t)$ of energy E and root-mean-square (Gabor) bandwidth β is transmitted at power P over bandwidth B. From the same observations the receiver must (i) estimate a sensing parameter—here the propagation delay τ, equivalently range $r = c\tau$—and (ii) decode data.

Sensing side (Fisher information → CRB). For delay estimation in additive white Gaussian noise, the Fisher information is the classical Gabor result $J(\tau) = 8\pi^2\beta^2(E/N_0)$, so every unbiased estimator obeys the Cramér–Rao bound $\mathrm{var}(\hat{\tau}) \geq 1/[8\pi^2\beta^2(E/N_0)]$, i.e. $\mathrm{var}(\hat{r}) \geq c^2/[8\pi^2\beta^2(E/N_0)]$ [36], [37]. This is exactly the per-link ranging floor against which GNSS front-ends have been designed since the 1980s.

GNSS specialization (the DOP decomposition). With n anchors and the linearized geometry matrix H (n rows; four columns, for three coordinates plus a clock term), and per-link variance σ_r^2 equal to the bound above, weighted least squares gives the position–clock covariance $C = \sigma_r^2\,(H^TH)^{-1}$. The scalar σ_r^2 is fixed by the waveform (β^2, E/N_0); the matrix $(H^TH)^{-1}$ is fixed by geometry, and its trace is the squared geometric dilution of precision, $\mathrm{GDOP}^2 = \mathrm{tr}[(H^TH)^{-1}]$. Thus positioning variance = (waveform-limited CRB) × (geometry-limited DOP), the separation every GNSS receiver reports [14], [48]. This is the Fisher-information matrix of the ISAC sensing problem, written in geodesy's native coordinates.

Coupling to the rate. The same signal carries information at the Shannon rate $R \leq B\log_2(1 + \rho)$, with SNR $\rho = P/(N_0B)$. Sensing and communication are coupled because they draw on one power, one bandwidth, and one aperture: with P fixed, shifting energy across B raises β^2 (tightening the delay CRB) while lowering ρ (reducing the per-Hz rate). The achievable pairs therefore trace a frontier in the (rate, sensing-precision) plane; the classical GNSS operating point—maximize β^2 for ranging, treat data as a low-rate side channel—is one extreme, and a pure data link the other. The full multi-user, multi-target, fading characterization is the open problem; the closed-form AWGN instance, with convex Pareto frontier and explicit optimal power split, is derived in the companion letter [25]. The chain—Fisher information ⇒ CRB ⇒ DOP-decomposed covariance ⇒ rate–distortion coupling through shared (P, B)—is followable here, without that letter (Figure 2).

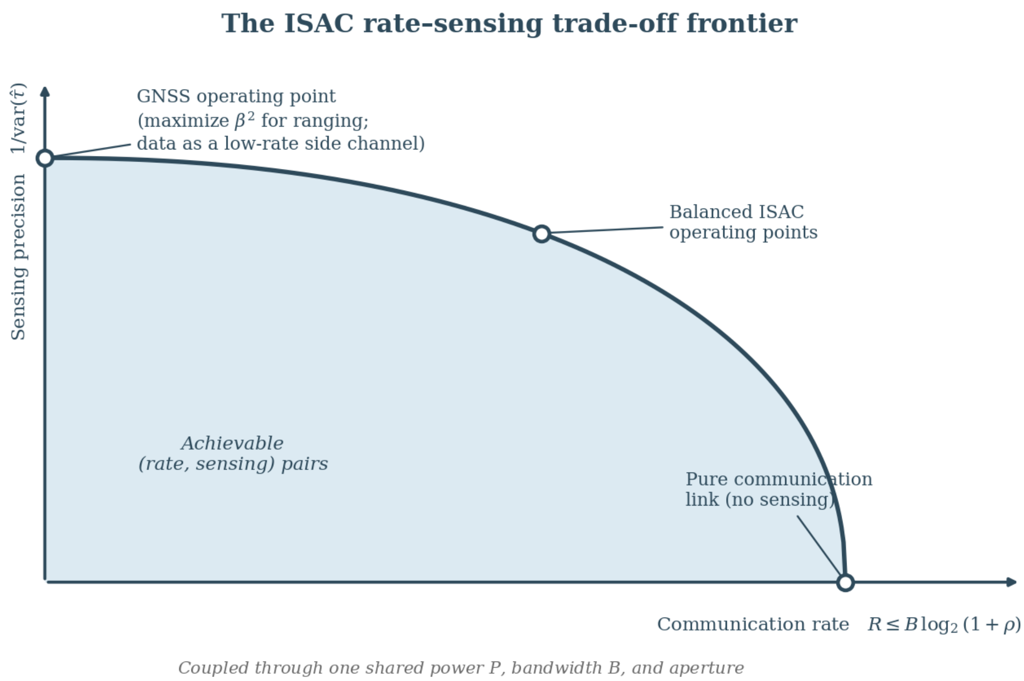

Figure 2. The ISAC rate–sensing trade-off frontier (illustrative). For a single shared waveform, the sensing-side Cramér–Rao bound and the communication-side Shannon rate are coupled through one power, bandwidth, and aperture budget. The classical GNSS operating point—bandwidth maximized for ranging, data carried as a low-rate side channel—is one extreme of the achievable region; a pure communication link is the other. The schematic shows the qualitative shape of the trade-off, not measured values; the closed-form additive-white-Gaussian-noise frontier is derived in [25].

4.2 Bridge 2: The Ambiguity Function and Delay-Doppler Channel Theory

The ambiguity function $\chi(\tau,\nu)$ of a signal $s(t)$ is the correlation between $s(t)$ and its delayed and Doppler-shifted copy $s(t-\tau)e^{j2\pi\nu t}$. For GNSS signal design the ambiguity function is the principal tool [14]—an object that originates in the radar tradition [36], [39]–[41]: the main-lobe width in τ governs ranging precision, the main-lobe width in ν governs tolerance to receiver dynamics, and the sidelobe structure governs robustness to multipath and cross-correlation between satellites. The Gold and Weil codes selected for GPS L1 C/A, the BOC and AltBOC modulations selected for Galileo E1 and E5, and the memory-code designs of BeiDou B1C are all the product of decades of ambiguity-function engineering. The delay-Doppler domain on which orthogonal time-frequency-space (OTFS) modulation is built [11] is the same domain: an OTFS symbol in the delay-Doppler plane, read through a time-varying multipath channel with a given delay spread and Doppler spread, convolves with the channel's delay-Doppler impulse response to produce the received signal. The channel's delay-Doppler kernel is, under mild assumptions, the ambiguity function of the channel impulse response evaluated against the transmitted symbol; the waveform-design problem—choose a pulse whose ambiguity function has a narrow main lobe and low sidelobes over the channel's support—is the problem GNSS signal designers have been solving since the 1970s. The Zak-transform formulation recently adopted in the Zak-OTFS literature is mathematically elegant and operationally equivalent to the quasi-periodic ambiguity-function formulation long used in radar and navigation.

The gap here is not conceptual but sociological. The delay-Doppler-domain capacity theorem—a Shannon-type result that establishes the information-theoretic limit of reliable communication over a time-varying channel described by its delay-Doppler kernel—remains unproven in general. The GNSS design intuitions that identify which kernel shapes are well-conditioned, which are pathological, and which admit closed-form receiver structures are the technical input such a proof requires. A capacity theorem derived without those intuitions risks being correct but unusable; with them, the theorem can be stated in a form that waveform designers recognize.

4.3 Bridge 3: Multi-Constellation Estimation and SAGIN Capacity

A modern GNSS receiver tracking signals from GPS, Galileo, BeiDou, and GLONASS simultaneously faces a fusion problem whose structure is independent of how the constellations are realized physically [14], [15]. Each constellation has its own reference frame, its own time scale, and its own set of broadcast biases; the receiver estimates a common user position and velocity along with per-constellation clock offsets known in the geodesy literature as inter-system biases. The optimal estimator is a weighted least-squares or Kalman fusion [46] whose weights depend on per-satellite signal quality, per-constellation ephemeris accuracy, and per-constellation atmospheric models. Robustness extensions—fault detection and exclusion, receiver autonomous integrity monitoring, multi-hypothesis ambiguity resolution—exist precisely because the fusion problem is non-trivial when constellations disagree.

Space-Air-Ground Integrated Networks (SAGIN) generalize this situation. A 6G or 7G user terminal may receive positioning-and-communication signals from low-Earth-orbit satellites, from high-altitude platforms, from terrestrial base stations, and from aerial relays simultaneously. Each layer has its own channel statistics, its own latency, its own reliability envelope, and its own bias structure. A unified capacity analysis for the SAGIN setting—the information-theoretic counterpart to multi-constellation GNSS positioning—is not yet available, and the reason is not a lack of modeling effort but a lack of the right fusion primitive. The geodetic multi-constellation estimator, recast in information-theoretic form, is that primitive: it handles heterogeneous measurement qualities, constellation-level bias parameters, and partial-observability gracefully, and it admits a closed-form convergence analysis that terrestrial-only heterogeneous-network capacity derivations typically lack. A concrete ML-ready instantiation of this fusion problem has been developed in the author's prior work on constellation-aware transformer architectures for multi-GNSS positioning [26], in which per-constellation embeddings and attention-weighted satellite selection reproduce the inter-system-bias accounting that classical receivers perform analytically while scaling to settings where the analytic weights are unavailable. We argue that the same architectural principle—treat each network layer as a constellation, each observation as a satellite—applies to SAGIN, though the formal analog awaits proof in a later phase of this program.

4.4 Bridge 4: Kalman Filtering and AI-Native Channel Tracking with Guarantees

The extended Kalman filter, as applied to tightly coupled GNSS–inertial integration, produces state estimates whose covariance is not only computed but is, under standard regularity conditions, a provable upper bound on actual estimation error [20]. This matters operationally: a Kalman-filtered position has an uncertainty ellipsoid that is meaningful, auditable, and the basis for integrity monitoring in safety-critical applications from aviation to autonomous vehicles. The filter's convergence properties, its sensitivity to model mismatch, and its failure modes under non-Gaussian disturbances have all been characterized over six decades of deployment [43], [44]. AI-native physical layers—air-interface designs in which core physical-layer functions such as channel estimation, equalization, and decoding are realized by models learned end-to-end rather than by fixed analytic algorithms—now propose to replace pieces of this pipeline with learned components—neural channel estimators, transformer-based equalizers, in-context-learning receivers that adapt on the fly to new channels without retraining. The empirical performance of these replacements is often excellent. The provable-guarantee analog of the Kalman covariance—an a-priori bound on the estimator's error distribution, whether stated as a convergence or regret bound, an integrity (error-covariance) bound, or a PAC-style generalization guarantee—is, for most of these architectures, not yet available. This is not a scholarly complaint. A wireless system whose estimator's error distribution cannot be bounded a priori is currently restricted from settings where integrity matters, and integrity matters in exactly the settings—aviation, autonomous mobility, medical telemetry—that 6G and 7G are proposing to serve.

Two strands of the author's prior work bear on this gap. The first is a body of results on in-context learning (ICL) complexity for navigation, geospatial, and climate tasks [27]–[33], which identifies the conditions under which a sufficiently large transformer can perform a given estimation task from in-context examples alone, and more importantly the conditions under which it cannot; these results specialize, to the navigation and geospatial setting, the broader theory of when transformers can implement standard estimation algorithms in context and select among them [52], [53]. The ICL-hardness results in that body of work are the natural mathematical analog of the observability conditions that govern Kalman-filter convergence. The second strand is a red-teaming critique of AI-safety evaluation methodology [34] whose central argument—that benchmark-driven validation systematically understates failure modes that matter in deployment—applies with particular force to AI-native physical layers that are currently evaluated on simulated channels with known statistics. The bridge this paper asserts is that Kalman-style provable guarantees and learned-component empirical performance are not alternatives between which 7G must choose. They are complements whose integration requires new theorems—specifically, convergence bounds for attention-based state estimators on specified function classes, and integrity-monitoring protocols for hybrid classical–learned filters.

4.5 Bridge 5: Antenna Phase-Center Calibration and Near-Field Electromagnetic Information Theory

A precision GNSS antenna does not radiate or receive from a point. Its effective phase center depends on elevation angle, azimuth, frequency, and—for choke-ring and other reduced-multipath designs—near-field interactions with the antenna's own structure. The geodetic community characterizes this behavior with phase-center-offset and phase-center-variation models, typically calibrated in anechoic chambers and distributed in standardized formats such as the International GNSS Service antenna calibration files [35]. A millimeter-level positioning application that ignored these models would produce millimeter-level bias; a kinematic application that ignored them would produce trajectory-dependent error that a Kalman filter cannot cleanly absorb.

Beyond-6G wireless is moving toward apertures—holographic surfaces, extremely large aperture arrays, reconfigurable intelligent surfaces of non-trivial size—at which the user terminal is demonstrably in the aperture's near field. In the near field, the plane-wave approximation that underlies classical MIMO information theory fails, and the capacity question becomes physically rather than abstractly infinite-dimensional. Electromagnetic information theory (EIT), an emerging sub-field, seeks to characterize the capacity available from such an aperture under Maxwell's equations and realistic radiation constraints [50], [51]. The parallel to antenna phase-center calibration is that, in both cases, a spatially extended radiator's behavior cannot be reduced to a single effective point without incurring bias, and that bias must be characterized before the information-theoretic question is meaningful. A concrete shared object already links the two sides: the spatial degrees of freedom of a field over a bounded aperture—finite and computable from geometry, and equal to the effective spatial bandwidth times the observation extent [50], [51]—are the electromagnetic counterpart of the number of independently resolvable parameters in a phase-center model, and the phase-center-variation map is precisely the per-direction bias that any degrees-of-freedom count must absorb before it can be read as a channel capacity. This bridge is more speculative than the first four. The phase-center-variation models are empirical fits, not closed-form physics, and the near-field capacity problem is not yet well enough posed for the mapping to be made rigorous. The claim here is not that a theorem is available but that a research direction is: near-field electromagnetic information theory in the presence of calibrated but non-ideal apertures is a problem on which geodetic antenna modeling has operational data that electromagnetic information theorists do not.

Collectively the five bridges describe what the next decade of foundational 7G theory work could look like if the geodetic and wireless communities engaged each other systematically. Whether that engagement happens on a timescale that matters is the subject of Section 5.

5. The Window Is Now

The argument made so far—that beyond-6G wireless has a theorem-shaped hole and that the geodetic-engineering tradition fills a significant fraction of it—does not by itself imply urgency. Urgency comes from the calendar. Three institutional processes now in motion will, by approximately 2030, convert what is currently a theoretical opportunity into a set of standards-bound engineering constraints; theoretical frameworks developed after those constraints are set will have to contort themselves to fit, while frameworks developed before can shape what the constraints become. The first is the ITU-R IMT-2030 framework [4]: the visioning phase concluded with publication of the framework recommendation, the performance-requirements and evaluation-methodology phase is in its final year at the time of writing, the radio-interface-technology submission and evaluation phase runs from 2027 through approximately 2029, and the IMT-2030 specifications are expected to be finalized in 2030. A foundational theorem published in 2026 or 2027 can inform what radio-interface proposals are submitted; a theorem published in 2031 cannot. The second is the World Radiocommunication Conference cycle: WRC-23 identified candidate bands in the 7–8 GHz, 14–15 GHz, and upper mid-band regions for possible future IMT use, with final allocations expected at WRC-27. Spectrum decisions are effectively irreversible on decade timescales, and the theoretical frameworks that inform band-specific design choices—near-field operation in upper bands, atmospheric absorption management, coexistence with passive services—are being selected now. Geodetic experience with GNSS bands in L-band and S-band, with ionospheric dispersion, with tropospheric delay, and with adjacent-band coexistence [22]–[24] is directly relevant to the upper-band decisions, but that experience has not been systematically injected into the WRC-27 preparation cycle. The third is 3GPP's standards timeline: Release 20, currently specifying the first 6G study items, runs through the second half of 2026, Release 21 is expected to contain the first normative 6G specification and runs from 2027 through mid-2028, and subsequent releases will elaborate 6G and eventually begin to look beyond it. The engineering community that authors these specifications looks to the research literature for conceptual tools; tools that exist as published theorems before Release 21 freezes will shape what Release 21 and its successors standardize. Figure 3 shows the 2026–2030 critical window for foundational 7G theory.

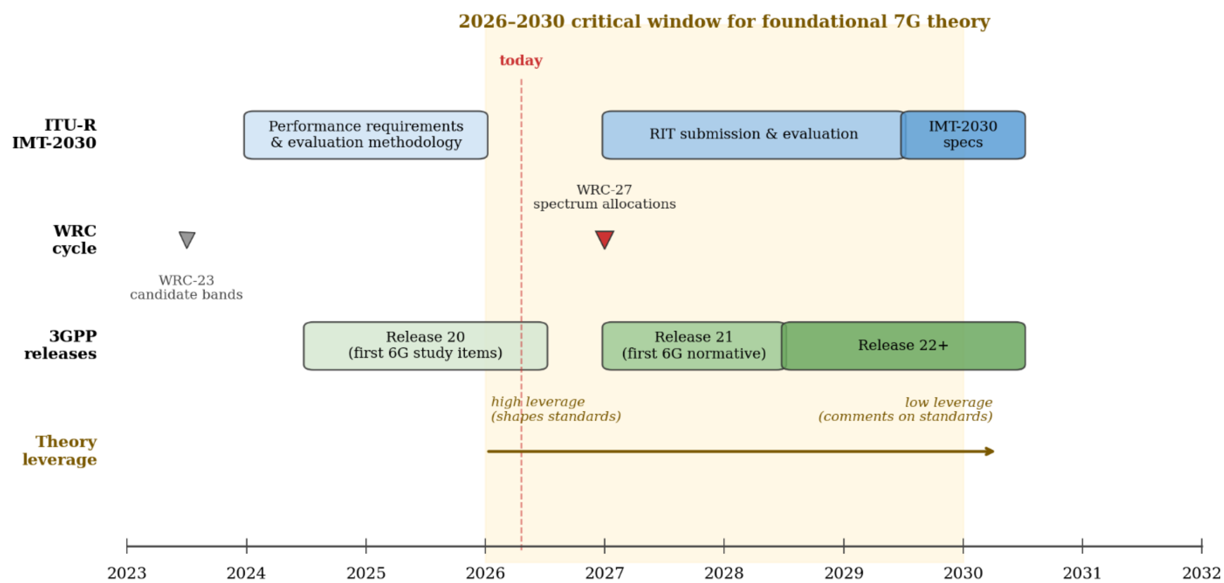

Figure 3. The 2026–2030 critical window for foundational 7G theory. Three institutional tracks set the calendar by which theoretical frameworks must be in place to influence what gets standardized: the ITU-R IMT-2030 evaluation phases, the World Radiocommunication Conference cycle (WRC-23 candidate-band identification, WRC-27 final allocations), and 3GPP's Release 20–22+ timeline. Theoretical leverage decays across the window: a result published before Release 21 freezes can shape what Release 21 normatively specifies; a result published after 2030 comments on what others have already specified.

The consequence is that the window for foundational work that actually constrains what 7G becomes is the period from now through approximately 2030. Work published after that window will still be valuable as academic contribution, but its ability to influence what systems get built will be substantially reduced. This is the present paper's principal pragmatic claim. It is a claim about leverage, not about quality, and it applies with equal force to the five bridges articulated in Section 4: each bridge, developed now, can shape a standards generation; each bridge, developed in 2032, will be commenting on one.

6. Limitations and Open Assumptions

The argument above is a position, not a proof, and several of its load-bearing assumptions deserve to be stated plainly.

Regime transfer is a hypothesis, not a result. The geodetic apparatus invoked here was developed for a regime that differs from beyond-6G ISAC in nearly every operating parameter: GNSS works at high effective SNR after long coherent integration, over narrow bands, with slowly varying and well-known satellite geometry and comparatively benign multipath. Terrestrial and non-terrestrial beyond-6G links will often operate at lower SNR, over wide bandwidths, under rapidly varying geometry, dense and time-selective multipath, and extreme user density. The bridges assert that the mathematical objects transfer—Fisher information, ambiguity functions, recursive estimators, fusion rules—not that GNSS numerical regimes or design heuristics transfer unchanged. Establishing the conditions under which each mapping survives the change of regime is part of the work this paper calls for, not a result it claims.

Radar and sonar share the same tools. The Fisher-information and ambiguity-function tradition behind Bridges 1 and 2 is not unique to geodesy; it originates in radar [39] and is equally central to sonar [42]. The claim is therefore not that geodesy alone possesses these tools, but that it has carried them furthest in the directions beyond-6G now needs: integer-ambiguity estimation at scale, multi-constellation fusion with inter-system biases, and deployed integrity monitoring under safety-of-life requirements. Where radar or sonar offer the more mature instrument—clutter modeling, waveform diversity—those traditions should be drawn on equally. The abstract's superlative should accordingly be read in a restricted sense: most mature for the positioning-integrity and multi-source-fusion sub-problems, not across all of sensing.

A single defining theorem may not exist. Section 2's pattern rests on three completed generational transitions, a small base even granting that it is the entire population of completed transitions. It remains possible that 6G and 7G are inherently multi-theoretic, with no single governing result, in which case the right object is a small family of theorems rather than "6G's theorem." Nothing in the five-bridges program depends on the single-theorem framing; the bridges stand or fall individually.

The five bridges are a selection, not a closure. They were chosen because each mapping can be stated crisply and the author has direct experience with the geodetic side; they are not claimed to be exhaustive, or even necessarily the five most important links. Bridge 1 is supported here by the self-contained sketch of Box 1 and, in full, by a companion letter [25] that is at present under review; until that review resolves, the "operational" label for Bridge 1 should be read as provisional. Bridges 2–4 are concrete but not yet proven, and Bridge 5 is explicitly a research direction rather than a result.

Program concentration. This paper is the perspective statement of a single-author research program, and a corresponding share of its citations are to that program's own recent and forthcoming work. Those self-references should be weighted accordingly: the independent literature cited throughout—in estimation and detection theory, radar and sonar, optimal filtering, multi-sensor fusion, information theory, and electromagnetic information theory—is the load-bearing support for the bridges, while the program's own papers are offered as worked instances rather than independent corroboration.

7. Conclusions: The Next Shannon May Come From an Unexpected Field

The call this paper makes is specific and symmetric. To the geodetic and GNSS research community: the integrated sensing and communication literature is where the mathematical tools you have developed for half a century are being rediscovered, often under different names and sometimes with weaker guarantees. The venues that publish this work—IEEE Transactions on Signal Processing, IEEE Transactions on Wireless Communications, IEEE Journal on Selected Areas in Communications, IEEE Transactions on Information Theory—are open to contributions that bring geodetic rigor to beyond-6G problems. The ION GNSS+ and IUGG communities have both the tools and the deployment experience that the wireless community needs. The barrier is bibliographic, not intellectual.

To the wireless-communications research community: the geodetic literature on Cramér–Rao bounds in the presence of multipath, on Kalman-filter convergence under model mismatch, on multi-constellation fusion with inter-system biases, on ambiguity-function design under dynamic conditions, and on antenna phase-center characterization in the near field contains worked solutions to problems that currently appear in beyond-6G papers as open questions. The journals that publish this work—Journal of Geodesy, Journal of Navigation, NAVIGATION: Journal of the Institute of Navigation, GPS Solutions—are a short reading list rather than a long one. The investment required to cross-train a graduate student in both traditions is a matter of two or three courses, not a second degree.

The pattern identified in Section 2—that every mature wireless generation was preceded by a foundational theorem—does not guarantee that the pattern will repeat for 7G. It is possible that seventh generation wireless will be defined, as some current treatments implicitly assume, as an incremental elaboration of 6G with no new foundational content. This paper has argued that the more interesting and more likely outcome is different: that 7G will be shaped by theorems not yet proven, that those theorems will draw on mathematical traditions currently outside the mainstream wireless literature, and that, for the positioning-integrity and multi-source-fusion problems at the core of ISAC, the geodetic tradition is the most mature and most immediately applicable of those traditions (Section 6 delimits this claim). The opening provocation asked what 6G's theorem is. The answer offered here is that 6G's theorem, and 7G's, may be written by someone who is currently not attending the conferences where 6G is discussed. The task of the present research program—of which this paper is the perspective statement, with a concurrent letter on the CRB–Rate tradeoff [25] and a program of foundational theorem papers to follow—is to ensure that those theorems get written in time to matter.

The next Shannon may come from an unexpected field.

8. Acknowledgments

This is an Artificial Intelligence-Assisted Research (AIAR). AI tools were employed as efficiency-enhancing assistants. The author declares no conflict of interest. This study received no external funding.

9. References

1. C. E. Shannon, “A mathematical theory of communication,” *Bell Syst. Tech. J.*, vol. 27, no. 3, pp. 379–423, Jul. 1948.
2. İ. E. Telatar, “Capacity of multi-antenna Gaussian channels,” *Eur. Trans. Telecommun.*, vol. 10, no. 6, pp. 585–595, Nov.–Dec. 1999.
3. E. Arıkan, “Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memoryless channels,” *IEEE Trans. Inf. Theory*, vol. 55, no. 7, pp. 3051–3073, Jul. 2009.
4. ITU-R, “Framework and overall objectives of the future development of IMT for 2030 and beyond,” Recommendation ITU-R M.2160, Nov. 2023.
5. I. F. Akyildiz, A. Kak, and S. Nie, “6G and beyond: The future of wireless communications systems,” *IEEE Access*, vol. 8, pp. 133995–134030, 2020.
6. H. Wymeersch *et al.*, “Cross-layer integrated sensing and communication: A joint industrial and academic perspective,” arXiv:2505.10933, 2025.
7. Y. Xiong, F. Liu, Y. Cui, W. Yuan, T. X. Han, and G. Caire, “On the fundamental tradeoff of integrated sensing and communications under Gaussian channels,” *IEEE Trans. Inf. Theory*, vol. 69, no. 9, pp. 5723–5751, Sep. 2023.
8. F. Liu, Y. Cui, C. Masouros, J. Xu, T. X. Han, Y. C. Eldar, and S. Buzzi, “Integrated sensing and communications: Toward dual-functional wireless networks for 6G and beyond,” *IEEE J. Sel. Areas Commun.*, vol. 40, no. 6, pp. 1728–1767, Jun. 2022.
9. H. Hua, J. Xu, and T. X. Han, “Optimal transmit beamforming for integrated sensing and communication,” *IEEE Trans. Veh. Technol.*, vol. 72, no. 8, pp. 10588–10603, Aug. 2023.
10. Y. Xiong, F. Liu, K. Wan, W. Yuan, Y. Cui, and G. Caire, “From torch to projector: Fundamental tradeoff of integrated sensing and communications,” *IEEE BITS Inf. Theory Mag.*, vol. 4, no. 1, pp. 73–90, May 2024.
11. R. Hadani, S. Rakib, M. Tsatsanis, A. Monk, A. J. Goldsmith, A. F. Molisch, and R. Calderbank, “Orthogonal time frequency space modulation,” in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, San Francisco, CA, USA, Mar. 2017, pp. 1–6.
12. V. Chamola, M. S. Peelam, M. Guizani, and D. Niyato, “Future of connectivity: A comprehensive review of innovations and challenges in 7G smart networks,” *IEEE Open J. Commun. Soc.*, vol. 6, pp. 3555–3613, 2025.
13. Z. Xing, “AIICN: A 7G multi-path transmission based on information-centric network,” *Front. Comput. Sci.*, vol. 7, Art. no. 1659616, Dec. 2025, doi: 10.3389/fcomp.2025.1659616.
14. E. D. Kaplan and C. J. Hegarty, Eds., *Understanding GPS/GNSS: Principles and Applications*, 3rd ed. Norwood, MA, USA: Artech House, 2017.
15. P. J. G. Teunissen and O. Montenbruck, Eds., *Springer Handbook of Global Navigation Satellite Systems*. Cham, Switzerland: Springer, 2017.
16. D. Medina, J. Vilà-Valls, A. Hesselbarth, R. Ziebold, and J. García, “On the recursive joint position and attitude determination in multi-antenna GNSS platforms,” *Remote Sens.*, vol. 12, no. 12, Art. no. 1955, 2020.
17. D. Das, L. Ortega, J. Vilà-Valls, F. Vincent, E. Chaumette, and L. Davain, “A new compact CRB for delay, Doppler and phase estimation—application to GNSS SPP and

RTK performance characterisation," *IET Radar, Sonar & Navigation*, vol. 14, no. 10, pp. 1537–1549, 2020.
18. F. Vincent, E. Chaumette, C. Charbonnieras, J. Israël, M. Aubourg, L. Ries, and F. Barbiero, "Asymptotically efficient GNSS trilateration," *Signal Process.*, vol. 133, pp. 270–277, Apr. 2017.
19. D. Medina, J. Vilà-Valls, E. S. Lohan, and A. Hein, "Performance limits of GNSS code-based precise positioning: GPS, Galileo and meta-signals," *Sensors*, vol. 20, no. 8, Art. no. 2196, Apr. 2020.
20. P. D. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*, 2nd ed. Norwood, MA, USA: Artech House, 2013.
21. P. J. G. Teunissen, "The least-squares ambiguity decorrelation adjustment: A method for fast GPS integer ambiguity estimation," *J. Geod.*, vol. 70, no. 1–2, pp. 65–82, Nov. 1995.
22. J. A. Klobuchar, "Ionospheric time-delay algorithm for single-frequency GPS users," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-23, no. 3, pp. 325–331, May 1987.
23. M. Hawarey and T. Ayan, "Global positioning system detection of Minuteman II launch and positioning of launch site," *J. Surveying Eng.*, vol. 131, no. 3, pp. 78–86, Aug. 2005, doi: 10.1061/(ASCE)0733-9453(2005)131:3(78).
24. M. Hawarey, "Travelling ionospheric disturbance over California mid 2000," *Nonlin. Processes in Geophys.*, vol. 13, pp. 1–7, 2006, doi: 10.5194/npg-13-1-2006.
25. M. Hawarey, "On the fundamental CRB–Rate tradeoff for GNSS-integrated sensing and communication," *AIR J. Eng. Technol.*, 2026, Art. no. AIRJET2026793, doi: 10.65737/AIRJET2026793.
26. M. Hawarey, "A constellation-aware transformer architecture for multi-GNSS positioning: Learned inter-system bias estimation and attention-based satellite selection," *AIR J. Eng. Technol.*, 2026, Art. no. AIRJET2026613, doi: 10.65737/AIRJET2026613.
27. M. Hawarey, "In-context learning characterization of navigation and GNSS foundation models: A theoretical framework for safety-critical positioning," *AIR J. Math. Comput. Sci.*, 2026, Art. no. AIRMCS2026525, doi: 10.65737/AIRMCS2026525.
28. M. Hawarey, "Characterizing in-context learning: When can transformers match standard learning algorithms?" *AIR J. Math. Comput. Sci.*, 2026, Art. no. AIRMCS2026277, doi: 10.65737/AIRMCS2026277.
29. M. Hawarey, "Quantum-enhanced in-context learning for geopotential field estimation: A theoretical framework," *AIR J. Math. Comput. Sci.*, 2026, Art. no. AIRMCS2026322, doi: 10.65737/AIRMCS2026322.
30. M. Hawarey, "ICL characterization of geo-foundation models," *AIR J. Math. Comput. Sci.*, 2026, Art. no. AIRMCS2026360, doi: 10.65737/AIRMCS2026360.
31. M. Hawarey, "ICL characterization of climate foundation models: When can transformers learn weather and climate?" *AIR J. Math. Comput. Sci.*, 2026, Art. no. AIRMCS2026405, doi: 10.65737/AIRMCS2026405.
32. M. Hawarey, "ICL characterization of multi-modal geo-foundation models: When can vision-language transformers learn geospatial tasks?" *AIR J. Math. Comput. Sci.*, 2026, Art. no. AIRMCS2026446, doi: 10.65737/AIRMCS2026446.
33. M. Hawarey, "Chain-of-thought ICL for GeoAI: Breaking the hardness barrier," *AIR J. Math. Comput. Sci.*, 2026, Art. no. AIRMCS2026482, doi: 10.65737/AIRMCS2026482.
34. M. Hawarey, "A red-teaming critique of AI-safety evaluation methodology," *AIR J. Interdiscip. Res.*, 2026, Art. no. AIRJIR2026369, doi: 10.65737/AIRJIR2026369.

35. R. Schmid, P. Steigenberger, G. Gendt, M. Ge, and M. Rothacher, “Generation of a consistent absolute phase-center correction model for GPS receiver and satellite antennas,” *J. Geod.*, vol. 81, no. 12, pp. 781–798, Dec. 2007.
36. H. L. Van Trees, *Detection, Estimation, and Modulation Theory, Part I*. New York, NY, USA: Wiley, 1968.
37. S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ, USA: Prentice Hall, 1993.
38. C. W. Helstrom, *Statistical Theory of Signal Detection*, 2nd ed. Oxford, U.K.: Pergamon Press, 1968.
39. P. M. Woodward, *Probability and Information Theory, with Applications to Radar*. London, U.K.: Pergamon Press, 1953.
40. M. I. Skolnik, *Introduction to Radar Systems*, 2nd ed. New York, NY, USA: McGraw-Hill, 1980.
41. N. Levanon, *Radar Principles*. New York, NY, USA: Wiley, 1988.
42. W. S. Burdic, *Underwater Acoustic System Analysis*. Englewood Cliffs, NJ, USA: Prentice Hall, 1984.
43. A. H. Jazwinski, *Stochastic Processes and Filtering Theory*. New York, NY, USA: Academic Press, 1970.
44. B. D. O. Anderson and J. B. Moore, *Optimal Filtering*. Englewood Cliffs, NJ, USA: Prentice Hall, 1979.
45. A. Gelb, Ed., *Applied Optimal Estimation*. Cambridge, MA, USA: MIT Press, 1974.
46. Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. Boston, MA, USA: Academic Press, 1988.
47. B. W. Parkinson and J. J. Spilker, Jr., Eds., *Global Positioning System: Theory and Applications*, vols. I–II. Washington, DC, USA: AIAA, 1996.
48. G. Strang and K. Borre, *Linear Algebra, Geodesy, and GPS*. Wellesley, MA, USA: Wellesley-Cambridge Press, 1997.
49. T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY, USA: Wiley, 1991.
50. O. M. Bucci and G. Franceschetti, “On the degrees of freedom of scattered fields,” *IEEE Trans. Antennas Propag.*, vol. 37, no. 7, pp. 918–926, Jul. 1989.
51. D. Dardari, “Communicating with large intelligent surfaces: Fundamental limits and models,” *IEEE J. Sel. Areas Commun.*, vol. 38, no. 11, pp. 2526–2537, Nov. 2020.
52. S. Garg, D. Tsipras, P. Liang, and G. Valiant, “What can transformers learn in-context? A case study of simple function classes,” in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 35, 2022, pp. 30583–30598.
53. Y. Bai, F. Chen, H. Wang, C. Xiong, and S. Mei, “Transformers as statisticians: Provable in-context learning with in-context algorithm selection,” in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 36, 2023.