AIR Journal of Engineering and Technology, Vol. 2026, ISSN: 3143-7494
Article Unique ID: AIRJET2026867
DOI: 10.65737/AIRJET2026867

Toward a Mathematical Theory of Integrated Positioning, Sensing, and Communication: Foundations for 7G

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

Received: 2026 August 13 | Revised: 2026 August 27 | Approved: 2026 August 31 | Published: 2026 September 04

Abstract

The sixth generation of cellular wireless has reached the normative standardization phase without an accompanying foundational theorem of the kind that preceded every prior generation: Shannon's coding theorem for 1G–3G, Telatar's multi-antenna capacity for 4G, Arıkan's channel polarization for 5G. Seventh-generation wireless is currently being defined at the level of use cases and marketing vocabulary; the mathematical problems that will actually constrain what 7G can become have not been posed. This survey argues that the foundational work required is substantial, that its natural language is the estimation-theoretic tradition developed over six decades of geodetic engineering, and that the open problems are numerous, well-defined, and presently unassigned. We organize the landscape around five research pillars; ISAC information theory, delay-Doppler capacity theory, Space-Air-Ground Integrated Network capacity with positioning constraints, AI-native physical-layer theory with provable guarantees, and electromagnetic information theory for near-field and holographic systems, and review the state of each pillar against its open problems. We argue that the Global Navigation Satellite System is the operational precedent for every system that Pillars I through V attempt to construct, and that the Cramér–Rao bound, Kalman filter, ambiguity function, dilution-of-precision formalism, multi-constellation estimation machinery, and antenna calibration theory of geodesy provide the mathematical toolkit that the current 6G and beyond-6G literature has not yet assimilated. We formalize a Fisher–Shannon unification object that appears across all five pillars, and we map the cross-pillar convergence zones—ISAC with OTFS, SAGIN with electromagnetic information theory, AI-native with semantic compression, quantum-enhanced positioning, and reconfigurable intelligent surfaces in integrated sensing—that will produce the highest-impact foundational results of the coming decade. We publish a numbered catalog of over one hundred open problems, each classified by difficulty, prerequisite knowledge, expected research effort, and the number of follow-on papers the problem is expected to generate. We close with a decade research roadmap tied to the ITU-R IMT-2030 evaluation phase, World Radiocommunication Conference 2027 spectrum decisions, and the 3GPP Release 20 through Release 22 standardization cycle, and we identify the 2026–2030 window as the critical period in which foundational theoretical work can still shape which systems get built.

Keywords: 7G; beyond-6G; integrated sensing and communication (ISAC); Cramér–Rao bound; Fisher information; Global Navigation Satellite System (GNSS); Kalman filtering; delay-

Doppler channel; Zak-OTFS; Space-Air-Ground Integrated Network (SAGIN); AI-native physical layer; electromagnetic information theory; holographic MIMO; near-field communication; foundational research roadmap.

Citation: Hawarey, M. (2026). Toward a Mathematical Theory of Integrated Positioning, Sensing, and Communication: Foundations for 7G. *AIR Journal of Engineering and Technology*, Vol. 2026, AIRJET2026867. https://doi.org/10.65737/AIRJET2026867

1. Introduction

This section motivates the survey. §1.1 states the historical pattern under which each prior wireless generation has rested on mathematics in place before the generation was inevitable, and observes that the pattern is broken for the sixth and seventh generations. §1.2 diagnoses the break as a discipline-level mismatch between the mathematical toolkit the foundational work requires and the toolkit the wireless-communications research community is most deeply trained in. §1.3 states the timing argument — why the foundational work matters now rather than later. §1.4 fixes the survey's scope, its five contributions, and the organization of the remaining sections.

1.1 The pattern: theorems before systems

Every mature generation of cellular wireless has rested on mathematics that was in place before the generation itself was inevitable. Shannon's 1948 characterization of the capacity of a noisy channel [1] supplied the objective function that the analog-to-digital transition of first through third generation systems was built to approach. Telatar's multi-antenna capacity analysis of Gaussian channels (circulated 1995, published 1999) [2] demonstrated that capacity scaled linearly in the smaller of the transmit and receive antenna counts and thereby converted multiple-input multiple-output architecture from a clever engineering proposal into the defining feature of fourth generation systems. Arıkan's 2009 channel polarization construction [3] produced the first explicit capacity-achieving code family of practical decoding complexity, and polar codes are now the control-channel code of fifth generation systems [4]. In each case a single citable result — publishable, criticizable, teachable — preceded the system and constrained what the generation could become. Figure 1.1 sketches the pattern, including the approximate lead times from publication of the theorem to the normative standardization moment it shaped.

A qualification on the pattern. The pattern admits counterexamples, and they bound the claim rather than refute it. Several technologies central to deployed cellular systems arrived without a preceding foundational theorem of the kind described above: orthogonal frequency-division multiplexing was an engineering construction whose theoretical analysis largely followed its adoption; turbo codes were discovered empirically in 1993 and their behavior is still not fully characterized analytically; and the second generation of cellular wireless rested on no new capacity result at all. The honest statement of the pattern is therefore narrower than a universal law. It is that the generations whose defining feature was a change in what the physical layer optimizes — the analog-to-digital transition, the shift to spatial multiplexing, the adoption of capacity-achieving codes — were each preceded by a result that specified the new objective, whereas the generations whose defining feature was an engineering refinement of an existing objective were not. This distinction matters for the argument that follows, because the beyond-6G capabilities surveyed here are of the first kind: integrated sensing, native positioning, and AI-native estimation each change what the physical layer is being asked to optimize, rather than improving how well it optimizes rate. A reader who rejects the historical pattern entirely can still accept the remainder of this survey, since the open problems catalogued in §10 stand or fall on their own mathematical merits.

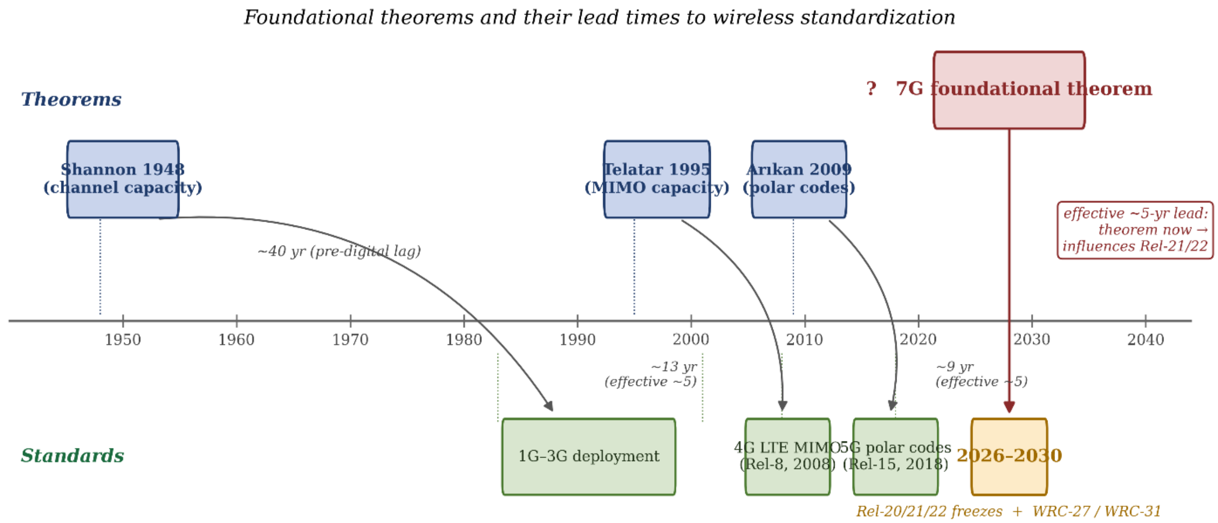

Figure 1.1. Historical-pattern timeline of foundational wireless theorems and their lead times to normative standardization. Shannon 1948 anchored 1G–3G digital-wireless adoption at an aggregate lag of approximately forty years (the pre-digital-technology lag is uninformatively long). Telatar 1995 anchored 4G LTE MIMO (3GPP Release 8, 2008) at about thirteen years' total lag, with an effective lag of roughly five years from result maturity to normative specification. Arıkan 2009 anchored 5G polar codes (3GPP Release 15, 2018) at about nine years' total lag, again roughly five effective years. The 2026–2030 window marks the 3GPP Release 20, 21, 22 freezes and the WRC-27 / WRC-31 spectrum cycles. A foundational theorem for 7G published during 2026–2028 would enter the standards process at the same effective five-year lead time that Telatar's and Arıkan's results enjoyed.

Sixth generation wireless has no such result. The ITU-R IMT-2030 framework is published [5], World Radiocommunication Conference 2027 will issue candidate spectrum allocations [6], and 3GPP Release 21 is scheduled to contain the first normative 6G specification [7]. What is missing is the theorem. The integrated-sensing-and-communication capacity-distortion region, the delay-Doppler channel Shannon capacity, the provable-guarantee analog of the Kalman covariance matrix for AI-native physical-layer processing, the electromagnetic information-theoretic capacity of a holographic aperture with physical mutual coupling — each is the subject of an active research literature [8], [9], [10], [11], and none has been established in the generality needed to play the role that Shannon's theorem, Telatar's theorem, and Arıkan's construction played for earlier generations.

If sixth generation wireless lacks a defining theorem, seventh generation wireless lacks even the problem statement. The term “7G” currently circulates in two contexts. The first is speculative end-of-survey sections appended to sixth-generation visions: a paragraph or two forecasting holographic communication, pervasive intelligence, or terahertz coverage beyond what 6G is expected to deliver. The second is use-case-oriented pieces describing entrepreneurial or commercial applications that the authors have decided to call 7G for rhetorical reasons. Neither genre produces a mathematical problem. The foundational work required to constrain 7G — to determine what it must optimize, what it cannot achieve, and what waveforms, architectures, and estimation procedures it will require — has not begun.

A positive criterion. This survey therefore uses the term in a specific and deliberately restrictive sense. A system is seventh-generation, for present purposes, if its performance is governed by a joint objective in which communication rate and estimation accuracy are simultaneously binding — that is, if the relevant performance object is a region in a multi-dimensional space of rate and error covariance rather than a scalar capacity, and if operating points on that region cannot be reached by optimizing rate and then treating sensing or positioning as a downstream service. Sixth-generation systems, on this criterion, are rate-governed systems to which sensing and positioning capabilities have been added; a seventh-generation system is one whose air interface is designed against the joint region from the outset. The criterion is not a marketing boundary and does not depend on a spectrum allocation, a release number, or a data-rate threshold. It is a statement about which mathematical object the design is optimizing, and it is the criterion under which the five pillars of §§3 through 7 are selected.

1.2 The diagnosis: a discipline-level mismatch

The diagnosis advanced in this survey is that the mismatch between the foundational work required and the foundational work being done is not a matter of insufficient effort from the wireless-communications research community. The community is producing thousands of papers per year on 6G topics. The mismatch is that the mathematical toolkit needed to formulate and attack the relevant 7G problems is, to a substantial degree, not the toolkit in which the wireless-communications community is most deeply trained. It is the toolkit of geodetic engineering and of the estimation-theoretic tradition more broadly.

Three observations support this diagnosis. First, the technologies most commonly invoked as beyond-6G enablers — integrated sensing and communication, positioning as a native wireless function, space-air-ground integrated networking, AI-native physical layers with safety guarantees, and holographic or near-field antenna systems — share a common mathematical structure: they are all fundamentally joint estimation-and-communication problems rather than pure communication problems. Second, the mathematical object best suited to describe performance in such systems is not a single scalar capacity but a region in a multi-dimensional performance space where communication rate trades against sensing accuracy, positioning error covariance, and physical-channel degrees of freedom. Third, the natural machinery for analyzing such regions — Fisher information matrices, Cramér–Rao bounds, Kalman filters, ambiguity functions, dilution-of-precision decompositions, multi-constellation fusion rules, atmospheric delay models, antenna phase-center calibration theory — is the mathematical core of geodetic engineering, which has used this machinery operationally and theoretically since the 1960s [12], [13], [14].

The Global Navigation Satellite System is the operational precedent. It is the only large-scale wireless system in current global service that simultaneously conveys information, measures its own channel with sub-nanosecond precision, estimates receiver position and velocity from those measurements, provides timing to continental electrical grids, observes the ionosphere and troposphere as a byproduct, and does all of this with standard receivers costing a few dollars per unit. In the language of the wireless-communications community, GNSS is the integrated-sensing-and-communication system, the positioning-and-communication system, the multi-constellation space-air-ground integrated network, and the low-complexity estimation-theoretic physical layer that 6G and 7G research programs have set themselves the goal of designing. The author's own research line, which originated in GPS signal processing and ionospheric detection of rocket launches using the Southern California GPS network [15], and has since extended to deep-learning architectures for multi-constellation positioning [16] and in-context learning characterization for navigation foundation models [17], sits at the center of this discipline-level mismatch and is the vantage from which the present survey is written.

1.3 Why this survey, and why now

The argument for publishing a foundational survey of beyond-6G mathematical problems at the present moment rests on three timing considerations.

First, the 2026–2030 window is when foundational theorems can still shape architecture. Standards bodies respond to mathematics when the mathematics arrives before the architecture is frozen. Telatar's multi-antenna capacity result circulated for approximately five years before 4G MIMO specifications stabilized [2], [18]. Arikan's polar-code construction had a similar lead time before 5G control-channel adoption [3], [4]. The ITU-R IMT-2030 evaluation phase runs through approximately 2029, with specifications expected in 2030; World Radiocommunication Conference 2027 will issue the candidate spectrum allocations that determine what 7G physics is available to build on; 3GPP Release 21 will freeze the first normative 6G air interface around 2028. Foundational theorems published after 2030 will inform textbook treatments but will arrive too late to constrain what the infrastructure becomes. The quantitative version of this timing argument is developed in Figure 1.1 above and is operationalized in the program-paper roadmap of §10.4.

Second, the open-problem catalog has become too large for any single research group to sweep. The survey below enumerates more than one hundred specific, well-defined foundational problems across five pillars — one hundred and twenty problems in the final catalog registered in §10. Each problem is of the size that a PhD student, postdoc, or small research group can address as a single paper. Many of the problems have been stated informally in the prior literature — appearing as “future work” in existing papers, as motivating remarks in survey sections, or as conjectures at the end of specific results. They have not been organized into a single list, classified by difficulty, or mapped against the prerequisite knowledge a researcher would need to attack them. The survey's catalog is meant to make that organization explicit.

Third, the cross-pollination between geodetic engineering and wireless communications has not happened by itself. Conferences are separate, journals are separate, departments are separate, and the overlap in the literature — where the mathematical identity of the problems is acknowledged — is thin. The present survey is, among other things, an argument for structured exchange: wireless-communications researchers reading the estimation-theoretic chapters of geodesy textbooks, geodetic engineers attending integrated-sensing-and-communication conferences, and individual researchers producing papers that cite both literatures explicitly. The companion position paper [19] makes this argument in shorter form and with a narrower focus on the five mathematical bridges that appear here as Pillars I through V.

1.4 Scope, contributions, and organization

This survey makes five contributions.

First, it presents the first comprehensive treatment that identifies seventh generation wireless as a distinct theoretical challenge — separate from 6G's remaining engineering agenda — and establishes the mathematical scope of that challenge at the level of specific open problems.

Second, it organizes the landscape around five research pillars — ISAC information theory, delay-Doppler capacity theory, SAGIN positioning-communication integration, AI-native physical-layer theory, and electromagnetic information theory — each with its own literature review, its own enumeration of open problems, and its own internal roadmap.

The principle of selection. The five pillars are selected by a single criterion, applied consistently: each is an area in which the performance object is a joint rate-and-estimation-accuracy region rather than a scalar capacity, and in which the estimation-theoretic machinery of geodetic engineering therefore applies directly. The criterion is deliberately exclusionary. Terahertz propagation, physical-layer security, energy efficiency, and semantic compression are all active and important beyond-6G research areas, and none of them is a pillar here, because in each case the foundational objects are scalar or single-discipline: terahertz work is dominated by propagation modeling and hardware constraint, security by equivocation-rate analysis, energy efficiency by scalar bits-per-joule optimization. Those areas appear in this survey where they intersect a pillar — semantic compression enters through Frontier 3 of §9.3, and reconfigurable surfaces through Frontier 5 of §9.5 — but they are not organized as pillars in their own right. The partition is accordingly not a claim to have enumerated the beyond-6G research agenda. It is a claim to have enumerated the portion of that agenda to which the geodetic estimation toolkit is the natural instrument.

Third, it formalizes the GNSS–wireless isomorphism: a precise mapping from the mathematical machinery of geodetic estimation theory to the currently open problems in each of the five pillars. The mapping is not a loose analogy. Dilution of precision is the right decomposition of the Fisher information matrix for any positioning-relevant ISAC design. The GNSS ambiguity function is the same Zak-transform object that Zak-OTFS waveforms exploit. Multi-constellation estimation is the right framework for SAGIN capacity-positioning integration. Kalman convergence is the right tool for analyzing recursive estimation in AI-native physical layers. Antenna phase-center-variation calibration is the mechanism by which electromagnetic information theory handles realistic aperture systems. Each mapping is established explicitly in the pillar where it applies, and the cross-pillar consolidation in §8 organizes them jointly via the Fisher–Shannon metric space and the concept-pillar matrix.

Fourth, it publishes a master catalog of one hundred and twenty numbered open problems, classified by difficulty (accessible, hard, grand challenge), by prerequisite knowledge, by estimated research effort in PhD-months, and by spiral potential (the expected number of follow-on papers each problem generates). The catalog is structured to be usable as a research-agenda-setting document for individual researchers, PhD programs, funding agencies, and research consortia.

Fifth, it presents a decade research roadmap that identifies Phase 0 flag-planting through Phase 4 institutionalization, aligned with ITU IMT-2030, WRC-27, and the 3GPP Release 20 through Release 22 standardization cycle, and identifies dependencies across pillars so that researchers can see which problems must be solved before which others become tractable.

The remainder of the survey is organized as follows. §2 reviews the state of sixth generation theory and establishes precisely where and how it falls short of the foundational standard that prior generations met. §§3 through 7 treat the five pillars in turn. §8 presents the Fisher–Shannon unification that operates across all five, together with the concept-pillar matrix that summarizes how the geodetic-estimation-theory toolkit distributes across the pillars. §9 maps the cross-pillar convergence frontiers where the highest-impact results of the coming decade will occur. §10 is the master open-problem catalog, with its 120-row classification table and the program-paper cross-reference. §11 is the decade roadmap, the call to the research community, and the conclusions. §12 states the limitations of the survey's argument, of its catalog metadata, and of the framework it proposes.

Scope and method of the literature survey. The literature reviewed here was assembled by searching IEEE Xplore, arXiv, and the 3GPP and ITU-R document repositories for the period January 2017 through April 2026, using the query terms integrated sensing and communication, ISAC, joint communication and sensing, delay-Doppler and OTFS, space-air-ground integrated network, LEO-PNT, AI-native air interface, electromagnetic information theory, and holographic MIMO, together with the corresponding geodetic terms Cramér–Rao bound, Fisher information matrix, dilution of precision, and carrier-phase ambiguity resolution. Foundational works predating 2017 were included where they are cited as anchors by the recent literature. A survey was classified as not attempting cross-pillar synthesis when its stated scope, section structure, and index of methods were confined to a single one of the five pillars defined in §1.4; the surveys so classified are those cited in §2.4. This procedure establishes what the present survey has examined, not what exists: the claim of §2.4 is that no cross-pillar synthesis was found among the surveys reviewed, and a systematic citation-graph bibliometrics of the geodesy–wireless boundary, which would test that claim properly, is itself registered as open problem O3.

1.5 Open problems from the introduction

The open problems registered in this module are numbered O1 through O4. Each problem is stated with a difficulty (Accessible, Hard, or Grand Challenge), a list of prerequisite areas of expertise, an estimated effort in PhD-months, and a spiral-potential estimate (the number of follow-on papers a solution is expected to generate). The complete catalog is consolidated in Section 10.

O1. The successor theorem. Establish, for the beyond-6G physical layer, a result that stands to the joint rate-and-estimation-accuracy regime as Shannon's capacity theorem stands to the digital transition, Telatar's multi-antenna capacity result to 4G, and Arıkan's polarization theorem to 5G: a provable characterization of the fundamental limit that the system architecture is thereafter designed to approach. Under the criterion of §1.3, such a result must govern a performance object in which rate and estimation accuracy are simultaneously binding, rather than a scalar capacity with an estimation-accuracy constraint appended. This is the problem toward which the remainder of the catalog is oriented; every other entry is either a component of it, a prerequisite for it, or a consequence of it. Difficulty: Grand Challenge. Prerequisites: information theory, estimation theory, convex analysis on the positive-semidefinite cone, and working familiarity with at least three of the five pillars. Effort: multi-year; not estimable in PhD-months at present. Spiral potential: defines a subfield.

O2. Responsiveness of the 2026–2030 standards window to foundational results. Determine empirically whether, and by what mechanism, a foundational theoretical result published within a standardization cycle influences the specifications produced by that cycle. The argument of §1.3 and §11.2 assumes such influence is possible and that the window closes; neither assumption is tested here. A study tracing the propagation of specific results into 3GPP and ITU-R specification text, with measured lead times, would establish whether the urgency claim of this survey is well founded. Difficulty: Hard. Prerequisites: standards-process analysis, bibliometrics, familiarity with 3GPP and ITU-R working procedures. Effort: 12–18 PhD-months. Spiral potential: informs research-programme design rather than generating follow-on theory papers.

O3. Geodesy–wireless citation-graph bibliometrics. Construct and analyse the citation graph linking the geodetic estimation literature to the beyond-6G wireless literature, and measure the extent of cross-citation across the boundary. This is the systematic test of the claim advanced in §2.4 and §5.1 — that the two literatures have developed the same estimation-theoretic machinery with minimal exchange. The present survey argues the claim from the surveys it examined, as §1.4 records; O3 is the study that would establish it. Difficulty: Accessible. Prerequisites: bibliometric methods, citation-graph analysis, domain familiarity sufficient to classify venues. Effort: 4–8 PhD-months. Spiral potential: 2–3 follow-on papers.

O4. A foundational-versus-architectural taxonomy for wireless results. Develop a principled taxonomy separating results that establish fundamental limits from results that specify architectures or procedures, and apply it to the beyond-6G literature. The distinction is used informally throughout this survey — it underlies the claim of §2.5 that engineering specification has outrun foundational theory — but it is not formally defined, and the boundary is contested in specific cases. Difficulty: Hard. Prerequisites: philosophy of applied mathematics, information theory, standards literature. Effort: 8–12 PhD-months. Spiral potential: 3–5 follow-on papers.

The target audience for this survey is threefold: wireless-communications researchers whose work could benefit from explicit engagement with estimation theory and geodetic engineering; geodetic engineers whose mathematical expertise applies directly to problems in the wireless-communications literature that those engineers may not have encountered; and PhD programs, funding agencies, and research consortia planning foundational programs for the decade ahead. No single subsection below is addressed exclusively to specialists in one of these groups; the survey is written to be read in full by all three.

2. State of 6G — Where Theory Stands and Where It Fails

This section provides the empirical ground for the survey's central claim: that sixth generation wireless has matured to the standardization phase without an accompanying foundational theorem, and that existing 6G surveys, though plentiful, do not address this lack. The section is organized in five subsections. §2.1 reviews the ITU-R IMT-2030 framework and extracts what it does and does not specify about foundational mathematics. §2.2 reviews 3GPP Release 18 through Release 20 and identifies where the gap between engineering work and foundational theory is already consequential. §2.3 surveys the vision and architecture 6G surveys and identifies their collective scope limitation. §2.4 surveys the technology-specific 6G surveys (ISAC, OTFS, SAGIN, holographic MIMO, AI-physical layer) and identifies the absence of cross-pillar synthesis. §2.5 consolidates these four reviews into the foundational-theorem-gap thesis and sets up the transition to the pillar framework of Sections 3 through 7.

2.1 The ITU-R IMT-2030 framework

The ITU-R M.2160-0 Recommendation [5] constitutes the globally agreed vision document for sixth generation wireless. It was approved by Radiocommunication Assembly 2023 in Dubai in November 2023, following completion at the 44th meeting of Working Party 5D in June 2023. The Recommendation identifies six usage scenarios—immersive communication, hyper-reliable and low-latency communication, massive communication, ubiquitous connectivity, artificial intelligence and communication, and integrated sensing and communication—and fifteen capability targets, of which nine extend capabilities present in IMT-2020 and six are new. The Recommendation is structured around four overarching design principles: sustainability, security and resilience, connecting the unconnected, and ubiquitous intelligence. In February 2026 Working Party 5D completed the draft follow-on Report on minimum requirements related to technical performance for IMT-2030 radio interface(s) [20], specifying the quantitative targets that candidate radio interface technologies must meet; that draft was submitted to Study Group 5 for approval scheduled December 2026, and was therefore not yet an issued Recommendation at the time of writing.

The IMT-2030 framework is, by design, not a theoretical document. It specifies what the system should do, not what the system is able to do under information-theoretic or estimation-theoretic constraints. Three of the six usage scenarios—artificial intelligence and communication, integrated sensing and communication, and ubiquitous connectivity (which in the framework's text includes non-terrestrial components)—invoke technical capabilities for which the foundational mathematical theory is incomplete. Fifteen capabilities are listed, but the values are stated as research and investigation targets with ranges rather than as results of capacity analyses, and the framework explicitly notes that “all values in the range have equal priority in research and investigation” [5]. This is the correct stance for a standards document. What it leaves uncovered is the question of whether any combination of the target values is simultaneously achievable. That question is the province of foundational theory, and the foundational theory has not been written.

A specific example makes the point concrete. The IMT-2030 framework identifies high-accuracy positioning (targeting centimeter-level accuracy in some deployment scenarios) and high data rates (peak-rate targets of 50–200 Gbps, depending on deployment scenario) as capability targets; ISAC is one of the six usage scenarios in which both capabilities are expected to coexist. The joint capacity-distortion region for an ISAC system that must simultaneously hit a centimeter-level positioning Cramér–Rao bound and a peak rate at the upper end of the IMT-2030 range on a shared waveform at a specified spectrum allocation has not been characterized in the general case. The closest published result is the CRB–Rate region of Xiong et al. under Gaussian channel assumptions [9], which is a foundational contribution but is neither sufficient to settle the IMT-2030 joint target nor yet translated into standards-usable bounds. The framework asks for the capability; the mathematics that says whether the capability is achievable has not been done.

2.2 3GPP Release 18 through Release 20

The Third Generation Partnership Project has, since its Release 18 work cycle (completed in 2024 with the "5G-Advanced" label), introduced artificial-intelligence and machine-learning capabilities at the physical layer through a series of study items [21]. The study items cover channel state information feedback compression with autoencoders, beam management with learned policies, and positioning with model-in-model approaches. Release 19 (in progress as of April 2026) extends these study items into normative specifications for specific use cases. Release 20 (Stage-3 freeze March 2027) is the final "5G-Advanced only" release: 3GPP has concluded that two releases are required to specify 6G, with Release 20 carrying the 6G study items and Release 21 carrying the normative work [22]. Release 21 (Stage-3 functional freeze December 2028, ASN.1/OpenAPI freeze March 2029) will contain the first full 6G air interface [7]. In parallel, ISAC has appeared as a study-item topic in Release 19 [23], with further work planned for Release 20 and 21.

The gap between engineering progress and foundational theory is, at this level, already operational. The AI/ML-for-PHY study items in Release 18 specify training datasets, model exchange formats, and performance-versus-overhead tradeoffs for specific use cases, but they do not rest on any provable bound that characterizes what a neural-network-based receiver can and cannot achieve as a function of model capacity, training-set size, and channel statistics. The ISAC study items in Release 19 describe sensing use cases and measurement procedures, but they do not rest on a capacity-distortion region that characterizes the joint sensing-communication performance region available under specified waveform and spectrum constraints. The absence of foundational theory has not prevented engineering progress. It has produced a body of engineering specifications whose performance claims are empirical rather than provable, and whose design decisions are consequently harder to defend, harder to extend, and harder to port to deployment scenarios outside the ones on which they were tested.

2.3 Vision and architecture surveys for 6G

The 6G survey literature includes a dense and valuable set of vision and architecture documents. Alsabah et al. [24] surveyed the broad technology enablers—terahertz communication, reconfigurable intelligent surfaces, non-terrestrial networks, artificial-intelligence-enabled networking—that 6G is expected to incorporate. Giordani et al. [25] provided an early systematic vision. Hexa-X, the European flagship 6G initiative, has produced the most extensive collective vision work, including a detailed account of the six research challenges (connecting intelligence, network of networks, sustainability, global service coverage, extreme experience, trustworthiness) [26] and topical reports on specific enablers. Nguyen et al. [27] surveyed 6G for the Internet of Things. Chamola et al. [28] produced a recent comprehensive landscape. The 2025 landscape-mapping survey [29] provides the most current systematic treatment of the framework.

These vision and architecture surveys are not foundational-theory surveys and are not intended to be. Their scope is the identification of use cases, the articulation of performance targets, and the enumeration of technology enablers. Within that scope they are thorough. The question they do not address—and do not attempt to address—is the question of which foundational theorems are required to constrain the system, which of those theorems have been proved, and which remain open. A reader of the vision and architecture surveys finds a picture of 6G that is architecturally complete but mathematically underspecified: one knows what the system is supposed to do, one knows what hardware and software components it is supposed to include, and one does not know what information-theoretic or estimation-theoretic bounds govern the achievability of the targets.

2.4 Technology-specific surveys

The 6G literature includes a second layer of surveys that treat specific technology pillars in depth. For integrated sensing and communication, Liu, Masouros, Eldar and collaborators [8] provided the definitive system-level treatment, followed by Zhang, Liu and collaborators' comprehensive "Intelligent ISAC" survey [30] and the Liu et al. fundamental-limits survey [31]. For the delay-Doppler channel family and OTFS/Zak-OTFS, the literature centers on the Hadani, Rakib, Molisch line from 2017 onward [32] and the Mohammed, Chockalingam, Calderbank Zak-OTFS formulation [33]. For Space-Air-Ground Integrated Networks, the major surveys are Liu, Shi and collaborators [34], Cheng, He, Yin and collaborators [35], and the recent Chen, Guo, Meng, Han, Li, Quek treatment of joint-communication-and-computing SAGIN [36]. For AI-native physical-layer theory, the surveys are more industry-led: Samsung's 6G AI/ML for Physical-Layer series [37]. the Science China Information Sciences overview by Cui et al. [38], and the Frontiers comprehensive review of AI-native 6G [39]. For electromagnetic information theory and holographic MIMO, the literature centers on Migliore's foundational work [40], the Pizzo, Marzetta, Sanguinetti Fourier plane-wave expansion [41], the two-part holographic MIMO tutorial of An, Yuen, Huang, Debbah, Poor and Hanzo [10], [42], and by Wang, Yang, Huang, and collaborators [11].

Each of these technology-specific surveys covers its pillar thoroughly. Collectively, they have two characteristics relevant to the present survey.

First, none of them attempts cross-pillar synthesis. The ISAC surveys do not treat the delay-Doppler channel as a native waveform for ISAC beyond the few papers that explicitly combine the two [33]. The SAGIN surveys do not treat positioning as a co-equal performance metric with communication rate; positioning is a use case to be supported, not a capacity-theoretic variable to be optimized against. The AI-native surveys do not engage seriously with the estimation-theoretic machinery that would be needed to provide provable guarantees for neural receivers in the presence of sensing constraints. The electromagnetic-information-theory surveys do not treat the near-field positioning problem as a direct target for EIT capacity results, despite the fact that the Fisher-information machinery for near-field positioning and the degree-of-freedom machinery for near-field capacity share mathematical structure. Each pillar is deep; the synthesis across pillars is absent.

Second, none of them frames its pillar as part of a specifically seventh-generation problem set. The surveys are written as 6G surveys. They treat their topics as 6G technologies. The language of a distinct post-6G generation requiring qualitatively new foundational mathematics—the language of this survey—does not appear. This is appropriate for the purposes the technology-specific surveys serve, but it leaves unoccupied the intellectual space that the present survey is meant to fill.

The companion position paper of the present research program [19] makes the case for the 7G framing in short form; the present survey develops it in full.

2.5 The foundational-theorem-gap thesis

The material above supports the following thesis.

Thesis. As of April 2026, sixth generation wireless has entered the normative standardization phase without an accompanying foundational theorem that plays the role Shannon's capacity result, Telatar's multi-antenna capacity characterization, or Arıkan's channel polarization construction played for earlier generations. The 6G survey literature, though extensive and valuable, does not collectively address this lack because its division of labor—architecture and vision surveys for breadth, technology-specific surveys for depth—does not include a category that treats the foundational-theorem question directly. The capability targets in the ITU-R IMT-2030 framework are specified as research targets because the underlying mathematics has not been done. The AI/ML and ISAC study items in 3GPP Release 18 and 19 are specified as engineering procedures because the foundational bounds that would constrain those procedures have not been established. The specific foundational problems whose solutions would fill this gap can be stated precisely, organized taxonomically, and attacked systematically. The survey that follows does so.

Three qualifications to this thesis are worth stating at the outset.

First, the thesis is not that no mathematical work is being done on beyond-6G topics. The Xiong et al. CRB–Rate region [9], the Kobayashi et al. capacity-distortion framework [43], the Ahmadipour et al. broadcast-channel ISAC results [44], the Lampel et al. DZT-OTFS capacity-equivalence result [45], the Pizzo-Marzetta-Sanguinetti electromagnetic-degree-of-freedom line [41], and the Wang et al. electromagnetic-information-theory synthesis [11] are all substantial foundational contributions. The thesis is that these contributions, taken together, do not yet constitute the foundational layer that a 6G-class (let alone 7G-class) wireless system requires, and that they have not been organized into a coherent research program with an explicit open-problem catalog.

Second, the thesis is not that geodetic engineering provides a complete solution. It provides a toolkit that is unusually well matched to the foundational problems of beyond-6G wireless, and it provides a body of operational precedent (the Global Navigation Satellite System) that demonstrates what an integrated sensing-positioning-communication system looks like when it works at scale. The toolkit and the precedent together form the starting point for the foundational work the present survey maps; they do not constitute the finished work.

Third, the thesis is not that the gap will be closed by any single research program, including the present one. The gap is large enough to occupy a substantial fraction of the 6G and 7G foundational research community for the next decade. The present survey's catalog of more than one hundred open problems is explicitly intended to be attacked in parallel by many researchers and groups, with cross-pollination across the pillars producing the highest-impact results.

With these qualifications in place, the remainder of the survey treats the five pillars in turn. Section 3 opens with Pillar I: ISAC Information Theory.

2.6 Open problems from the state of 6G theory

The open problems registered in this module are numbered O5 through O9, following the conventions stated in §1.5.

O5. From IMT-2030 capability targets to information-theoretic bounds. Derive, for each of the fifteen IMT-2030 capability targets of Recommendation ITU-R M.2160 [5], the corresponding information-theoretic or estimation-theoretic bound, and determine which targets are provably achievable, which are provably unachievable under stated channel assumptions, and which remain open. The Recommendation states targets as research objectives; none is presented as the consequence of a capacity analysis. Difficulty: Hard. Prerequisites: information theory, estimation theory, propagation modelling, familiarity with the IMT-2030 framework. Effort: 18–24 PhD-months. Spiral potential: 6–10 follow-on papers, one or more per capability target.

O6. A provable-bound framework for AI/ML physical-layer procedures. Establish a framework within which the 3GPP Release 18 and Release 19 AI/ML physical-layer study items [21], [22] can be assigned provable performance bounds rather than empirical benchmark results. The procedures are being specified; the bounds do not exist. This problem is the state-of-6G-level statement of the void developed in detail in Pillar IV. Difficulty: Hard. Prerequisites: statistical learning theory, information theory, 3GPP physical-layer procedures. Effort: 18–30 PhD-months. Spiral potential: 5–8 follow-on papers.

O7. A usage-scenario-to-pillar taxonomy. Map the six IMT-2030 usage scenarios and fifteen capability targets onto the mathematical structures required to analyse them, and determine which pillars of §§3–7 — or which structures outside that partition — each scenario requires. The partition of this survey is one such mapping, constructed from the selection principle of §1.4; a systematic mapping constructed from the Recommendation outward would test it. Difficulty: Accessible. Prerequisites: familiarity with ITU-R M.2160, information theory, estimation theory. Effort: 6–12 PhD-months. Spiral potential: 3–5 follow-on papers.

O8. Empirical lead times of foundational theorems. Measure, for the foundational results of prior generations — Shannon 1948, Telatar 1995/99, Arıkan 2009, and candidate counterexamples such as OFDM and turbo codes — the interval between publication and appearance in normative specification, together with the mechanism of transmission. The three-case induction of §1.1 and the timing argument of §11.2 both rest on lead times that are asserted from the historical record rather than systematically measured. Difficulty: Accessible. Prerequisites: bibliometrics, standards archaeology, telecommunications history. Effort: 4–8 PhD-months. Spiral potential: 2–3 follow-on papers.

O9. A framework for identifying foundational gaps prospectively. Develop a method for identifying, within a live standardization cycle, the foundational results a generation lacks — rather than identifying them retrospectively once the generation has been deployed. This is the general form of the objection stated in §12.1: that foundational results are identifiable as such only in hindsight, and that the gap this survey identifies may be a gap in recognition rather than in mathematics. A prospective method would answer the objection or establish that it cannot be answered. Difficulty: Hard. Prerequisites: history and philosophy of engineering science, information theory, standards-process analysis. Effort: 12–18 PhD-months. Spiral potential: 3–5 follow-on papers.

3. Pillar I — ISAC Information Theory

Integrated sensing and communication is the pillar at which the discipline-level mismatch identified in §1 becomes mathematically explicit. The wireless-communications literature has, since approximately 2018, invested considerable analytic effort in characterizing the joint capacity-distortion region of a channel that simultaneously conveys information and supports an estimation task on the same waveform. The Global Navigation Satellite System community has, since approximately 1960, operated exactly such a system at global scale, with a mature mathematical apparatus for the estimation side—Cramér–Rao bounds on pseudorange and range-rate, dilution-of-precision decompositions of the Fisher information matrix, Kalman filtering of multi-epoch measurements, ambiguity-function analysis of ranging codes, multi-constellation fusion under systematic-error constraints—that has essentially no analog in the ISAC information-theory literature. This section surveys the ISAC information-theory literature, introduces the GNSS–ISAC isomorphism as a formal object, and defines the Fisher–Shannon metric space that will organize the survey's subsequent treatment of the other four pillars. Pillar I is the pillar at which the survey's central claim—that the foundational work beyond-6G wireless requires is expressible in the mathematical language of geodetic estimation theory—acquires its cleanest and most detailed justification.

3.1 The ISAC information-theory literature

The ISAC information-theory literature in its modern form begins with Kobayashi, Caire and Kramer [43], who introduced the capacity-distortion framework for a state-dependent memoryless channel in which the transmitter conveys a message while the receiver estimates a state parameter associated with the channel. The Kobayashi-Caire-Kramer formulation treats sensing as a state-estimation task, specifies a distortion metric on the estimated state, and characterizes the achievable set of (rate, distortion) pairs through a single-letter optimization.

The framework is general enough to accommodate arbitrary distortion measures and arbitrary state distributions. Its generality is also a limitation: the achievable-region boundary is specified implicitly by an optimization that must be solved channel-by-channel, and closed-form solutions exist only for special cases. The Kobayashi-Caire-Kramer paper nonetheless fixed the vocabulary—“capacity-distortion region,” “state-dependent channel with generalized feedback,” “sensing as state estimation”—that subsequent ISAC information theory has used.

Ahmadipour, Kobayashi, Wigger and Caire [44] extended the capacity-distortion framework to the broadcast channel, treating a two-user setting in which a single transmitter serves two independent receivers while simultaneously estimating state information. The broadcast-channel extension introduces the complication that the sensing task couples to the communication task through the rate-splitting and binning constructions by which broadcast-channel capacity regions are characterized, and the achievable regions generally admit no simpler representation than the general inner-bound expressions. The significance of the broadcast extension is that it demonstrates that the capacity-distortion framework is not a peculiarity of the point-to-point setting; it scales, in principle, to the multi-user settings that practical wireless systems actually operate in.

The most influential recent result in ISAC information theory is that of Xiong, Liu, Cui, Yuan, Han and Caire [9], who characterized the fundamental tradeoff of ISAC under Gaussian channels. Their analysis identifies two distinct tradeoff phenomena. The first, which they term the “Subspace Tradeoff,” describes how the transmit covariance matrix must allocate its eigen-energy between the subspaces favorable for communication and those favorable for sensing. The second, the “Deterministic-Random Tradeoff,” describes how the waveform's statistical properties must compromise between the deterministic-pilot behavior favorable for parameter estimation and the Gaussian-random behavior favorable for achieving Shannon capacity. The Subspace-and-Deterministic-Random framework produces a complete characterization of the CRB–Rate region for MIMO Gaussian ISAC channels and is the reference result against which subsequent work is now routinely benchmarked. The deterministic-random tradeoff was subsequently developed in a rate-distortion perspective by Liu, Xiong, Wan, Han and Caire [46], and extended by Xiong, Liu and Lops [47] to characterize the sensing-optimal operating point. The Xiong-Liu line is, in the terminology of this survey, the closest existing analog to a foundational theorem for ISAC information theory.

Parallel to the Xiong-Liu line, Li, Andrei, Djuhera, Mönich and Boche [48] have developed capacity-distortion tradeoffs in memoryless ISAC systems with explicit attention to the structural properties of the achievable region under specific distortion measures. Dong, Liu, Xiong, Cui and collaborators [49] have extended ISAC information theory to networked settings with sensory-data sharing across nodes, opening the question of distributed ISAC capacity. The “Intelligent ISAC” survey by Zhang, Liu and collaborators [30] consolidates the state of the art as of late 2025 and identifies a number of the open problems that appear in the subcatalog at the end of this section. The Peng and Wu treatment of joint source-channel coding for ISAC [50] addresses the separation-theorem question: under what conditions can the source coding of the sensing task and the channel coding of the communication message be performed independently without loss of optimality? They formulate a joint source-channel coding framework for a state-dependent memoryless channel in which the transmitter estimates the channel state from generalized feedback while conveying a source sequence under a communication-distortion constraint, characterize the resulting capacity-distortion-cost tradeoff, and prove that separate source and channel coding attains joint optimality in that setting. The result is therefore a positive separation theorem for the single-user memoryless case. What remains open is whether separation survives the extensions that matter for the positioning-relevant problems of this survey: multi-user topologies, extended targets, and channels with memory.

The survey of fundamental limits by Liu and collaborators [31] is the most comprehensive recent catalog of ISAC fundamental-limit results. It is the natural entry point for a reader coming to ISAC information theory from outside and it identifies, with the appropriate caveats, where the existing results are tight, where they are loose, and where no result is yet available. [51] provides a complementary monostatic-bistatic overview. The JSAC treatment by Liu, Masouros, Eldar and collaborators [8] remains the standard system-level reference, and the Zhang-Liu intelligent-ISAC survey is its 2025 update.

3.2 Survey taxonomy for ISAC information theory

The ISAC information-theory literature admits a natural taxonomy along three axes, which the present survey adopts.

The first axis distinguishes the *sensing-objective formulation*. Capacity-distortion formulations (Kobayashi-Caire-Kramer, Ahmadipour et al., Li-Boche) treat sensing as the recovery of a state variable from the received signal under a user-specified distortion metric; the sensing performance is then a distortion figure computed from the posterior distribution of the state. CRB-rate formulations (Xiong et al., Liu et al., Dong et al.) treat sensing as parameter estimation and compute sensing performance via the Cramér–Rao bound on the estimation error. The two formulations are not equivalent in general. Distortion-based formulations are more flexible—any distortion metric can be accommodated—but produce implicit achievable-region characterizations. CRB-based formulations are more restrictive—they require a parameter-estimation setting with differentiable likelihood—but produce explicit characterizations in terms of Fisher information matrices. For the positioning-relevant ISAC problems that are this survey's principal interest, the CRB formulation is the natural one.

The second axis distinguishes the *network topology*. Point-to-point ISAC is the setting treated by Kobayashi-Caire-Kramer and by Xiong et al., in which a single transmitter-receiver pair communicates while simultaneously estimating channel state. Multiple-access ISAC treats a setting in which multiple transmitters contend on a shared channel while a single receiver communicates with each and estimates joint state; this setting is addressed partially in the Xiong et al. framework and more completely in the Dong et al. sensory-data-sharing line. Broadcast-channel ISAC is the Ahmadipour et al. setting in which a single transmitter serves multiple receivers while performing sensing; the capacity regions here are characterized only via inner bounds except in degenerate cases. Relay-channel ISAC, interference-channel ISAC, and networked ISAC (arbitrary graph topology) are either sparsely treated or not treated at all in the information-theoretic literature; they are among the open problems catalogued in §3.6.

The third axis distinguishes the *waveform structure*. Deterministic-waveform ISAC treats the sensing waveform as known to the receiver prior to the sensing observation; this is the setting of classical radar and of GNSS ranging, and it produces the best possible CRB for a given signal energy. Random-waveform ISAC treats the sensing waveform as a realization of a random process that is itself the communication-information-bearing object; this is the setting in which ISAC information theory engages capacity most directly, and it is the setting in which the Xiong-Liu "Deterministic-Random Tradeoff" produces its strongest conclusions. Semi-deterministic waveforms—where a portion of the waveform is deterministic (pilots, reference signals) and a portion is random (data)—are the practical compromise that most contemporary wireless systems adopt; their ISAC information-theoretic characterization is active research.

The three axes together produce a three-dimensional taxonomy within which each ISAC information-theoretic result can be placed. A single-cell entry in this taxonomy—point-to-point, CRB-rate, random-waveform, Gaussian channel—is the Xiong-Liu 2023 result. Most other cells are either open or partially open; the survey's subcatalog in §3.6 enumerates the open cells.

3.3 The GNSS–ISAC isomorphism

The Global Navigation Satellite System is, in its operational realization across GPS, GLONASS, Galileo, BeiDou, QZSS and IRNSS, a collection of broadcast transmitters (the satellites) and a population of receivers (the users) operating a simultaneous communication-and-sensing task: each satellite conveys navigation data (the communication message—ephemeris, clock corrections, ionospheric model parameters, almanac, satellite health) on a waveform that the receiver uses for pseudorange estimation (the sensing task) and from which the receiver jointly estimates position, velocity and time. The mathematical apparatus that the geodesy literature has developed for this task is extensive and is in exact mathematical correspondence with the apparatus that the ISAC information-theory literature is presently constructing. This subsection formalizes that correspondence as the "GNSS–ISAC isomorphism."

The ranging CRB as sensing distortion. The Cramér–Rao bound on pseudorange estimation for a GNSS signal of bandwidth B, received at carrier-to-noise ratio C/N_0, integrated over coherent integration interval T is, in its simplest form, $\sigma^2_\tau \geq 1 / (8\pi^2 \beta^2_{rms} (C/N_0) T)$, where β_{rms} is the root-mean-square bandwidth of the ranging signal [12], [13]. This bound is mathematically identical to the CRB on time-delay estimation that enters the Xiong et al. CRB–Rate region characterization when the sensing task is specialized to monostatic ranging and the waveform covariance is restricted to the code-division-multiple-access family used by GNSS. What the geodesy literature calls "the ranging-CRB-to-C/N0 relation" and the ISAC literature calls "the sensing-CRB branch of the CRB-rate region" are the same object, expressed in different notational traditions.

Dilution of precision as Fisher-information decomposition. The dilution-of-precision formalism of GNSS—in which the receiver-position covariance matrix resulting from pseudorange estimation across N satellites is the product of the ranging variance σ^2_τ and a geometry-dependent DOP factor computable from the N line-of-sight unit vectors—is, in the language of estimation theory, exactly the decomposition of the inverse Fisher information matrix for the position estimator into a signal-strength component and a geometric component

[14]. For an ISAC system that must estimate a target position from distributed sensing observations, the inverse Fisher information matrix admits the same decomposition. What the ISAC literature is presently deriving on a case-by-case basis for specific sensing geometries is, in the geodesy literature, a standard formalism that has been operationally deployed for six decades. The GNSS literature's DOP-on-the-map-as-a-function-of-time-of-day computations are running instantiations of the Fisher-information decomposition that ISAC information theory is rediscovering.

The P0-2 result as the first explicit bridge. The program's companion technical letter [52] makes this correspondence explicit for the single-input-single-output case. It derives a closed-form CRB–Rate region for a SISO ISAC system in which the sensing task is GNSS-style pseudorange estimation on a spread-spectrum ranging signal sharing the channel with data; it recovers the classical GNSS ranging-CRB expression as a limiting case at the pure-sensing corner of the region; and it recovers the geometric dilution-of-precision as the structure of the multi-satellite extension. The letter is the first formal result known to the present author in which a CRB–Rate region in the information-theoretic sense and the classical GNSS DOP formalism are shown to be different faces of the same mathematical object. For the restricted setting it treats — single-input single-output, Gaussian, one-way ranging — the correspondence is therefore established as a theorem rather than asserted as an analogy. It is important to be precise about the scope of that statement: what has been proved is a limiting-case recovery within one special case, not a structure-preserving bijection between the two theories in general. The term "isomorphism" is used throughout this survey in the weaker sense of a checkable dictionary between mathematical objects, and the extent to which the dictionary survives generalization to matrix-valued, multi-user and two-way-ranging settings is precisely the content of the open problems catalogued below. Figure 3.1 illustrates the structure of the CRB–Rate region with the GNSS DOP boundary case, the Shannon capacity boundary case, and the joint-ISAC interior operating point identified explicitly.

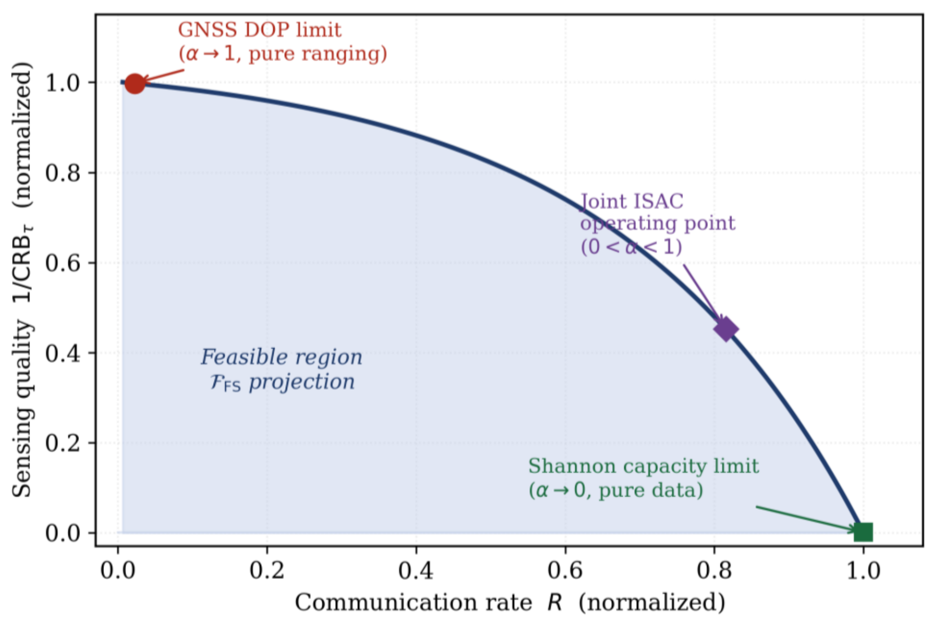

Figure 3.1. CRB–Rate region in the Fisher–Shannon metric space, projected onto (rate, 1/CRB_τ) axes. The Pareto boundary is convex; the upper-left corner (α → 1) recovers the classical GNSS ranging-CRB and dilution-of-

precision limit, the lower-right corner ($\alpha \to 0$) recovers the Shannon capacity of the data-only channel, and interior points correspond to joint ISAC operating regimes. The figure schematizes the closed-form region established by [52] in the SISO Gaussian setting; the matrix-valued boundary for the MIMO setting (problem O10 below) generalizes the same geometry.

Kalman filtering as the temporal completion. The geodesy literature's standard treatment of positioning does not stop at the single-epoch CRB. It uses the Kalman filter, with process noise covariance derived from the dynamics model (constant velocity, constant acceleration, Wiener process on velocity, etc.) and measurement covariance derived from the ranging CRB, to propagate the joint state estimate across measurement epochs [12]. The Kalman-filter framework provides a complete theory of the error-covariance dynamics, including the asymptotic steady-state covariance for recurring measurement geometries, the transient-response covariance following measurement gaps or outages, and the covariance response to measurement outliers. The corresponding temporal completion of ISAC information theory—a Kalman-filter analog in which the covariance that is being propagated is the error covariance of a joint communication-sensing estimate, with process noise on the target state and measurement noise characterized by the instantaneous CRB–Rate operating point—has not been constructed. The present program treats this construction in paper P1-4 (Kalman-ISAC Theory).

The ambiguity function as the common object. The Zak-domain radar ambiguity function of Woodward [53], the Doppler ambiguity function of classical pulse-Doppler radar, and the code-delay ambiguity function of GNSS ranging are, in the formulation of the modern Zak transform, the same object up to the choice of basis [33]. Pillar II treats the Zak-domain ambiguity function as the central object of delay-Doppler capacity theory; Pillar I treats it as the resolution-and-sensitivity structure of ISAC sensing; and the geodesy literature treats it as the autocorrelation structure of GNSS ranging codes. The three treatments are in mathematical correspondence. The reader who is entering this material from one community will recognize its structure from the other two.

Taken together, these correspondences establish what the survey refers to throughout as the *GNSS–ISAC isomorphism*: a specific, checkable mapping between the mathematical objects of geodetic estimation theory and the mathematical objects of ISAC information theory under which the results of each literature can be transcribed into the other. The isomorphism is not complete—it has well-defined limits, which the open problems in §3.6 and in subsequent pillars identify—but the portion of it that is complete is already sufficient to support the survey's central claim that the foundational work beyond-6G wireless requires is, at its mathematical core, an elaboration of the estimation-theoretic machinery that geodetic engineering has developed.

3.4 The Fisher–Shannon metric space

The correspondence established in §3.3 motivates a unifying mathematical object which the present survey names the Fisher–Shannon metric space—a working name whose metric structure is deferred, as §3.4 makes precise—and which is carried forward as the integrating object for all five pillars. This subsection introduces the object formally enough to make its role in the survey precise; the full theorems and proofs that characterize it are deferred to the program paper P1-1 (SISO Fisher–Shannon Capacity-CRB Region).

Object definition. For an ISAC setting with communication input distribution p_X, channel transition probability $p_{Y|X,\theta}$, and sensing parameter θ whose posterior the receiver estimates from the channel output Y, the Fisher–Shannon metric space $\mathscr{F}_{FS}$ is the set of pairs (Rate, FIM^{-1}) attainable as p_X ranges over admissible input distributions, where Rate is the mutual-information rate I(X; Y) and FIM^{-1} is the inverse Fisher information matrix for estimating θ from Y. The object is a region rather than an established metric space: equipping it with a metric requires a specific choice of distance on the cone of positive-semidefinite matrices, and the several Löwner-order-compatible candidates—among them the affine-invariant Riemannian distance and the log-Euclidean distance—are not equivalent. This survey therefore treats $\mathscr{F}_{FS}$ as the (Rate, FIM^{-1}) region under the Löwner partial order, with the rate component carrying the standard nat scale, and defers the choice of metric and the verification of its axioms to the program paper P1-1. Conjecture FS-2 below is stated with respect to the cone-product partial order induced by the Löwner order on the FIM^{-1} component, and does not depend on that choice.

Why this object, and why now. The CRB–Rate region of Xiong et al. is the two-dimensional projection of $\mathscr{F}_{FS}$ onto (rate, scalar-CRB) axes obtained by replacing the FIM with its trace or by specializing to scalar parameters. The capacity-distortion region of Kobayashi-Caire-Kramer is a different projection, onto (rate, scalar-distortion) axes obtained by a distortion-metric-induced scalarization. The dilution-of-precision formalism of geodesy is the geometric decomposition of points in $\mathscr{F}_{FS}$ into an SNR factor and a geometry factor. Treating the matrix object FIM^{-1} directly, rather than scalarizing it at the outset, is the move that unifies these three treatments and that connects them to the Kalman-filter framework (where FIM^{-1} is the instantaneous measurement-covariance input to the filter update), to the Zak-transform structure (where FIM^{-1} decomposes along the delay-Doppler axes), and to the AI-native PHY setting (where FIM^{-1} is the natural object against which learned-receiver performance should be benchmarked).

Five conjectures about $\mathscr{F}_{FS}$. The survey states the following five conjectures explicitly here. The conjectures are to be proved (or disproved) in P1-1 and its successors; their statement here is intended to orient the subsequent pillar sections, each of which will encounter $\mathscr{F}_{FS}$ in a specific specialization.

Conjecture FS-1 (Achievability-by-Gaussian-inputs): For the Gaussian ISAC channel with linear observation model and Gaussian prior on θ, the boundary of $\mathscr{F}_{FS}$ is achieved by Gaussian input distributions, and the boundary admits a closed-form parametric representation.

Conjecture FS-2 (Convexity): $\mathscr{F}_{FS}$, when viewed as a subset of $\mathbb{R}_{\geq 0} \times PSD_d$ with d the dimension of θ, is a convex set under the natural cone-product order.

Conjecture FS-3 (Multi-user tensorization): For a K-user MAC–ISAC setting with independent users, the K-user Fisher–Shannon metric space is the convex hull of tensor products of single-user regions, with appropriate corrections for shared-waveform-sensing gains.

Conjecture FS-4 (Kalman-limit recovery): The steady-state covariance achievable by a Kalman filter operating on measurements drawn from an ISAC system at Fisher–Shannon operating point (Rate, FIM^{-1}) admits a closed-form expression in terms of FIM^{-1} and the process-noise covariance.

Conjecture FS-5 (GNSS-DOP recovery): For multi-transmitter ISAC settings with geometry defined by N lines of sight to the sensing target, the FIM^{-1} projection onto position coordinates admits a geometry-factor decomposition that specializes to the classical GNSS DOP under the one-way-ranging restriction.

These five conjectures provide the scaffold against which Pillars II through V will be organized. Pillar II's open problems correspond to Fisher–Shannon characterizations along the delay-Doppler axes of the FIM. Pillar III's open problems correspond to Fisher–Shannon characterizations for multi-constellation SAGIN geometries. Pillar IV's open problems concern learned-receiver approximations to the $\mathscr{F}$_FS boundary. Pillar V's open problems concern the FIM^{-1} entries that come from near-field electromagnetic structure rather than from far-field scalar propagation.

3.5 Toward AI-native ISAC: the constellation-aware transformer

The preceding subsection framed the Fisher–Shannon metric space as the object that classical estimation theory constructs. The parallel object in the neural-network-based estimation literature is the posterior that a learned inference network approximates when trained on samples from the ISAC channel. The two objects are not the same in general: the classical FIM^{-1} is a lower bound on the error covariance achievable by any unbiased estimator, whereas the learned posterior is an achievable-error specification that a specific neural-network architecture attains under a specific training regimen. The gap between the two—the amount by which learned posteriors fall short of the CRB—is the quantity that AI-native physical-layer theory needs to characterize, and it will recur as the central open question of Pillar IV.

The Constellation-Aware Transformer architecture of the author [16] is an existence proof that transformer-style attention mechanisms can be adapted to the multi-constellation positioning problem. The architecture treats each satellite observation as a token, conditions attention on the satellite-system embedding (GPS, GLONASS, Galileo, BeiDou, QZSS, IRNSS) and on the per-satellite ephemeris information, and produces a posterior over receiver position from an arbitrary and variable set of simultaneously tracked satellites. The architecture has been shown to adapt to constellation-availability perturbations that break classical least-squares solvers—partial outages, ionospheric-event distortions, multipath-induced ranging biases—while retaining the Fisher-information-theoretic structure through the attention mechanism's implicit weighting of tokens by information content. The architecture is not, at present, accompanied by a capacity-theoretic bound: the question of whether the constellation-aware transformer, asymptotically in its training set and in its parameter count, approximates the classical FIM^{-1} of the underlying multi-constellation geometry is open.

This question admits a natural ISAC information-theoretic formulation. Given an ISAC channel operating at a specific Fisher–Shannon point (Rate, FIM^{-1}), and a neural-network receiver with P trainable parameters trained on N independent realizations of the channel, what is the error covariance achievable by the neural receiver as a function of P and N? The question is the ISAC specialization of the general PAC-learning question about sample complexity of capacity-achieving codes; it is less studied than the pure-communication version, and it is open in substantial generality. The survey treats this question as problem O19 in the subcatalog below.

3.6 Pillar I open-problem subcatalog

The open problems registered in this module are numbered O10 through O31. Each problem is stated with a difficulty (Accessible, Hard, or Grand Challenge), a list of prerequisite areas of expertise, an estimated effort in PhD-months, and a spiral-potential estimate (the number of follow-on papers a solution is expected to generate). The complete catalog is consolidated in Section 10.

O10. Full Fisher–Shannon boundary for MIMO Gaussian ISAC. Establish the complete boundary of the Fisher–Shannon metric space $\mathscr{F}$_FS for the MIMO Gaussian ISAC channel with Gaussian prior on the sensing parameter. The Xiong-Liu 2023 result supplies the (rate, scalar-CRB) projection; the full matrix-valued boundary remains open. Difficulty: Hard. Prerequisites: MIMO capacity theory, Fisher information, convex optimization on the PSD cone. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers on specific specializations (fading, sparse, multi-target).

O11. Proof of Conjecture FS-1. Prove or disprove that Gaussian inputs achieve the boundary of $\mathscr{F}$_FS for the Gaussian ISAC channel. Difficulty: Hard. Prerequisites: entropy-power inequality, moment-matching arguments for Gaussian optimality. Effort: 12–18 PhD-months. Spiral potential: 2–3 follow-on papers on non-Gaussian extensions.

O12. Proof of Conjecture FS-2. Prove or disprove convexity of $\mathscr{F}$_FS under the cone-product order. Difficulty: Hard. Prerequisites: convex analysis on PSD cones, information-geometric structure of Fisher information. Effort: 8–12 PhD-months. Spiral potential: 2 follow-on papers on boundary-achieving distributions.

O13. Proof of Conjecture FS-3 (MAC tensorization). Prove or disprove the tensorization property for K-user MAC–ISAC. Difficulty: Hard. Prerequisites: MAC capacity theory, distributed-estimation Fisher-information bounds. Effort: 12–18 PhD-months. Spiral potential: 3–5 follow-on papers on interference, broadcast, and relay extensions.

O14. Broadcast-channel ISAC boundary. Extend the Ahmadipour et al. inner bounds to a full characterization of the broadcast-channel ISAC capacity-distortion region for the Gaussian setting. Difficulty: Grand Challenge. Prerequisites: broadcast-channel capacity theory (Marton, superposition, rate-splitting), capacity-distortion framework. Effort: 24–36 PhD-months. Spiral potential: defines a sub-literature.

O15. Interference-channel ISAC. Characterize the achievable CRB-rate region for the two-user Gaussian interference channel with simultaneous sensing. Difficulty: Grand Challenge (open even in the pure-communication version). Prerequisites: interference-channel capacity, Han-Kobayashi coding, ISAC tradeoff characterization. Effort: 24–36 PhD-months. Spiral potential: defines a sub-literature.

O16. Relay-channel ISAC. Characterize the achievable CRB-rate region when the relay performs simultaneous communication-and-sensing. Difficulty: Hard. Prerequisites: relay-channel capacity, cooperative-sensing Fisher-information combination. Effort: 18–24 PhD-months. Spiral potential: 3–5 follow-on papers.

O17. Extended-target ISAC CRB. Extend the Xiong-Liu CRB-rate characterization from point-target sensing to extended-target sensing, where the target occupies a finite spatial extent and the sensing task is to estimate a target-shape-parameter vector. Difficulty: Hard. Prerequisites: extended-target radar theory, Fisher-information for shape parameters, MIMO ISAC. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers on specific shape models.

O18. Wideband ISAC capacity-CRB region. Extend the CRB-rate characterization from narrowband Gaussian to wideband (frequency-selective, dispersive) Gaussian ISAC channels. Difficulty: Hard. Prerequisites: wideband capacity theory, wideband-CRB for delay-spread channels. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers (OFDM-ISAC, single-carrier-ISAC, OTFS-ISAC).

O19. Neural-receiver ISAC sample complexity. Characterize the gap, as a function of neural-network parameter count and training-set size, between the error covariance achievable by a learned posterior and the Fisher–Shannon FIM^{-1} for a given Gaussian ISAC channel. Difficulty: Grand Challenge. Prerequisites: PAC-learning theory, neural-network approximation theory, Fisher information. Effort: 18–30 PhD-months. Spiral potential: defines a sub-literature at the Pillar I / Pillar IV interface.

O20. Continuous-waveform ISAC CRB-rate. Extend the CRB-rate characterization to continuous-time waveforms with bandwidth constraints, producing bounds expressed directly in signal-bandwidth and coherent-integration-interval rather than in discrete-time block length. Difficulty: Hard. Prerequisites: continuous-time capacity, continuous-time CRB for delay estimation. Effort: 12–18 PhD-months. Spiral potential: 2–3 follow-on papers bridging discrete-time ISAC and classical-radar continuous-time analysis.

O21. Kalman-ISAC steady-state covariance. Prove Conjecture FS-4: derive a closed-form expression for the steady-state Kalman-filter covariance on measurements drawn from an ISAC operating point. Difficulty: Hard. Prerequisites: Kalman-filter theory, Lyapunov-equation analysis, ISAC CRB-rate region. Effort: 12–18 PhD-months. Spiral potential: defines the Kalman-ISAC sub-literature (the program's P1-4 paper is a first entry).

O22. Kalman-ISAC transient response. Characterize the transient-covariance response of a Kalman filter on ISAC measurements following a measurement-geometry change (e.g., constellation reconfiguration, satellite rise/set, outage). Difficulty: Hard. Prerequisites: Kalman-filter transient analysis, ISAC CRB-rate. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers on specific geometry-change scenarios.

O23. Joint source-channel coding ISAC separation. Determine the conditions under which the source-coding (sensing) and channel-coding (communication) components of an ISAC system can be separated without capacity loss. The Peng-Wu 2026 result establishes separation for the single-user state-dependent memoryless channel; the multi-user, extended-target and memory-bearing cases are open, and it is not known whether separation is preserved when the sensing objective is a Fisher-information rather than a distortion functional. Difficulty: Hard. Prerequisites: joint source-channel coding theory, ISAC capacity-distortion. Effort: 12–18 PhD-months. Spiral potential: 2–3 follow-on papers on specific non-separable settings.

O24. Multi-target ISAC CRB-rate. Extend the point-target CRB-rate characterization to the multi-target setting with unknown target count, producing a joint detection-estimation-communication region. Difficulty: Hard. Prerequisites: multi-target tracking, random-finite-set filters, ISAC capacity. Effort: 18–24 PhD-months. Spiral potential: 4–5 follow-on papers.

O25. Privacy-constrained ISAC capacity. Characterize the capacity-distortion region of an ISAC channel subject to a privacy constraint on the sensing information (the sensing output must satisfy a differential-privacy bound with respect to user attributes). Difficulty: Hard. Prerequisites: differential privacy, information-theoretic privacy, ISAC CRB-rate. Effort: 12–18 PhD-months. Spiral potential: 2–3 follow-on papers on specific privacy models.

O26. Fading-channel ISAC capacity. Characterize the capacity-distortion region for ISAC under standard fading models (Rayleigh, Rician, Nakagami) in both ergodic and outage formulations. Difficulty: Accessible–Hard (depends on model). Prerequisites: fading-channel capacity, fading-channel estimation. Effort: 8–15 PhD-months per model. Spiral potential: 5–7 follow-on papers.

O27. Finite-blocklength ISAC. Develop the finite-blocklength (Polyanskiy-Poor-Verdú) refinement of the CRB-rate region, characterizing the second-order (dispersion) term. Difficulty: Hard. Prerequisites: finite-blocklength information theory, ISAC CRB-rate. Effort: 15–20 PhD-months. Spiral potential: 3 follow-on papers on URLLC-ISAC applications.

O28. ISAC under hardware impairments. Characterize how quantization, phase noise, I/Q imbalance and nonlinear power-amplifier distortion modify the ISAC capacity-distortion region. Difficulty: Accessible–Hard. Prerequisites: hardware-impaired capacity, impairment-aware estimation. Effort: 8–15 PhD-months. Spiral potential: 4–6 follow-on papers on specific impairment combinations.

O29. ISAC with imperfect CSI. Characterize the CRB-rate region when channel-state information is imperfect at the transmitter and/or receiver. Difficulty: Hard. Prerequisites: imperfect-CSI capacity, ISAC CRB-rate. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O30. ISAC energy efficiency. Characterize the capacity-distortion-energy region, adding energy consumption as a third performance axis. Difficulty: Hard. Prerequisites: energy-efficient capacity, ISAC CRB-rate. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers on green-ISAC.

O31. ISAC capacity under secrecy constraints. Characterize the capacity-distortion region subject to a secrecy constraint on the communication message against an eavesdropper who may also be performing sensing. Difficulty: Hard. Prerequisites: physical-layer secrecy, ISAC. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

Twenty-two problems are registered in this module, at the targeted density of 20–25. Together they cover the three axes of the ISAC-IT taxonomy (sensing-objective formulation, network topology, waveform structure), the temporal-completion direction (Kalman), the cross-pillar interfaces (Pillar II via wideband and OTFS, Pillar III via SAGIN geometries, Pillar IV via neural receivers), and the realistic-system directions (hardware impairments, imperfect CSI, secrecy, energy, privacy).

4. Pillar II — Delay-Doppler Capacity Theory

The second pillar of the survey is the delay-Doppler channel family and its emerging information theory. Delay-Doppler is the domain in which the channel's scattering structure—the collection of physical reflectors that produce the multipath realization seen by the receiver—appears sparsely and stably, in contrast with the rapidly varying representations produced in the time and frequency domains by the same underlying scattering. The argument for delay-Doppler waveforms rests on the observation that a channel which is doubly dispersive in time and frequency is, in the Zak-transform-induced delay-Doppler basis, a channel whose support is concentrated on a small set of delay-Doppler points that move only as the physical geometry changes. Orthogonal Time-Frequency Space (OTFS) modulation and its successor Zak-OTFS are the waveform families that exploit this concentration. The information-theoretic question that this pillar addresses is the status of capacity theorems native to the delay-Doppler representation: under what conditions does a Shannon-type capacity result exist for channels described as stochastic processes on the delay-Doppler plane, and what are those capacities? The current answer, set out below, is that a substantial equivalence result exists in the literature and is widely misunderstood as a native delay-Doppler capacity theorem; that a genuinely native theorem for sparse delay-Doppler channels has not been established; and that the ambiguity-function structure that ties delay-Doppler waveform design to GNSS ranging codes suggests a specific attack angle that the pillar develops.

4.1 From OFDM to OTFS to Zak-OTFS

Orthogonal frequency-division multiplexing, which remains the default cellular-wireless waveform from fourth through sixth generations, converts a frequency-selective channel into a collection of narrowband flat-fading subchannels and supports equalization by per-subcarrier scalar division. OFDM performs well when the channel is only weakly time-selective, so that the Doppler spread is a small fraction of the subcarrier spacing; it performs poorly when the Doppler spread is large, as in high-mobility vehicular channels, low-Earth-orbit satellite links, and millimeter-wave channels with moving blockers, because the Doppler-induced inter-carrier interference defeats the single-tap-per-subcarrier equalizer structure.

Orthogonal Time-Frequency Space modulation, introduced by Hadani, Rakib, Molisch and collaborators in a sequence of 2017–2018 papers [32], treats the information symbols as samples on the delay-Doppler plane rather than on the time-frequency plane. The OTFS transmitter applies a two-dimensional inverse symplectic Fourier transform to map the delay-Doppler symbol array into the time-frequency plane, passes the resulting signal through a multicarrier modulator, and transmits it. At the receiver, the inverse operation recovers the delay-Doppler symbols. The channel seen by the information symbols, in the absence of imperfections, is a two-dimensional convolution on the delay-Doppler plane with a kernel that is exactly the sparse delay-Doppler channel representation of the physical multipath. The property that gives OTFS its robustness to high Doppler is that the delay-Doppler kernel is sparse and slowly varying, so that equalization reduces to a two-dimensional convolution-inverse problem on a sparse support rather than a per-subcarrier division on a rapidly varying channel.

The original OTFS formulation used a particular pulse-shaping construction that was not fully compatible with the symplectic Fourier structure of the delay-Doppler plane. Mohammed, Hadani, Chockalingam and Calderbank, in a sequence of papers culminating in [33] and [54], reformulated the system using the Zak transform as the primary analytical object, producing "Zak-OTFS" as a mathematically cleaner version of the original idea. In the Zak-OTFS formulation, the information symbols are samples of a Zak-domain function; the transmitted signal is the inverse Zak transform of that function, suitably pulse-shaped; and the channel is represented as a Zak-domain operator whose structure reflects the physical delay-Doppler spreading. The Zak-OTFS formulation exposes a crystallization condition—a precise relationship between the symbol spacing in the delay-Doppler plane, the pulse-shape support, and the channel's delay-Doppler spread—under which the receiver sees exactly the intended sparse-kernel convolution without aliasing [54]. The crystallization condition is the Zak-domain analog of the Nyquist condition and is the mechanism by which Zak-OTFS achieves its best-case robustness properties. Figure 4.1 illustrates the two central objects of this pillar: the sparse-scatterer support of the physical channel on the delay-Doppler plane, and a representative Zak-domain pulse shape with its crystallization cell.

Subsequent work has extended Zak-OTFS in several directions. Mehrotra, Mattu and Calderbank [55] developed Zak-OTFS with spread-carrier waveforms that reduce peak-to-average power ratio while preserving the Zak-domain structure. Jayachandran et al. introduced interleaved-pilot designs that support channel estimation at the Zak-domain level [56]. Recent work has combined Zak-OTFS with integrated sensing, producing embedded-pilot schemes in which the Zak-domain reference signal serves simultaneously as communication-channel pilot and as sensing waveform. The IEEE Communications Society course on Zak-OTFS and delay-Doppler communications [57] indicates that the topic has entered the mainstream pedagogical literature.

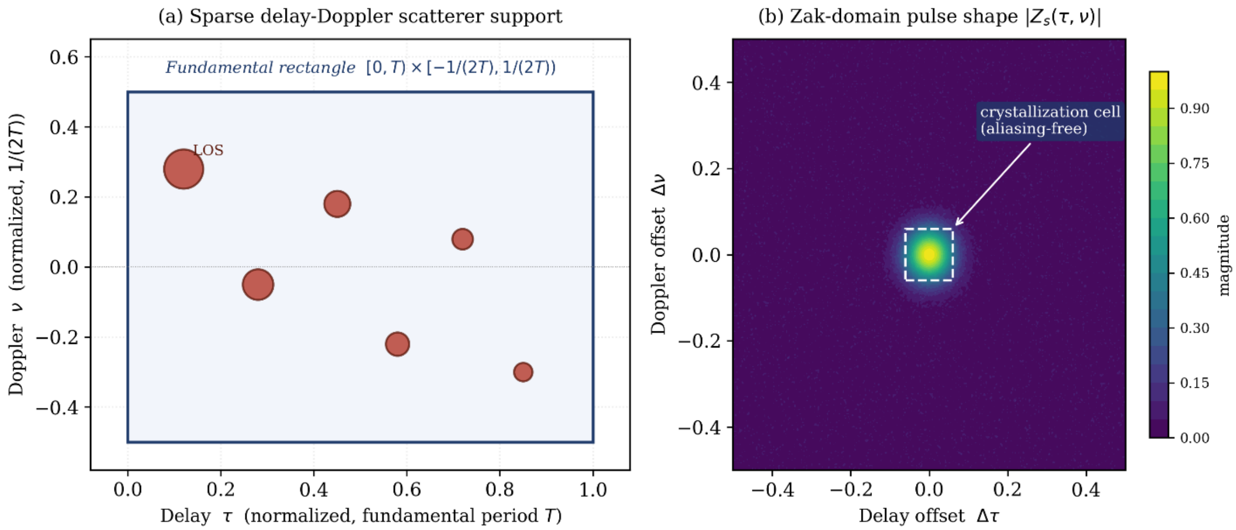

Figure 4.1. Delay-Doppler representation of a Zak-OTFS channel. (a) The physical multipath channel is sparse on the delay-Doppler plane: a small number of scatterers (physical reflectors) populate the fundamental rectangle [0, T) × [−1/(2T), 1/(2T)), with the largest amplitude typically corresponding to the line-of-sight (LOS) path. The locations move only as the physical geometry of the channel changes. (b) Zak-domain magnitude $|Z\_s(\tau,\nu)|$ of a thumbtack-like pulse shape; the dashed white box marks the crystallization cell, the region within which Zak-OTFS achieves aliasing-free reception under [54] condition. The two panels together show why Zak-OTFS is robust to high-Doppler channels: the sparse channel support of (a) interacts with the localized pulse of (b) through a two-dimensional convolution whose inversion is well-conditioned.

4.2 The Zak transform as the native delay-Doppler basis

The central analytic object of Pillar II is the Zak transform itself. For a continuous-time signal s(t) and a positive parameter T, the Zak transform is defined by $Z\_s(\tau, \nu) = \Sigma\_\{n \in \mathbb{Z}\} s(\tau + nT) e^\{-j2\pi\nu nT\}$ on the fundamental rectangle [0, T) × [−1/(2T), 1/(2T)), and it satisfies the quasi-periodicity relations that make it the appropriate tool for analyzing signals whose natural structure lies jointly in time (delay) and frequency (Doppler). The Zak transform has three properties relevant to the present discussion [58], [33].

First, the Zak-domain representation of a doubly dispersive channel is the channel's delay-Doppler spreading function, which for physical multipath channels is supported on a sparse set of delay-Doppler points corresponding to the underlying reflectors. The Zak domain is therefore the representation in which the physical channel structure is most compact.

Second, the Zak transform connects delay-Doppler analysis to the radar ambiguity function of Woodward [53]. For a signal s(t) with autocorrelation $R\_s(\tau, \nu) = \int s(t) s^*(t - \tau) e^\{-j2\pi\nu t\} dt$ (the radar ambiguity function), the Zak-transform magnitude $|Z\_s(\tau, \nu)|$ encodes the same information about the signal's delay-Doppler resolution as R_s encodes for radar applications. Woodward's uncertainty principle—that the volume under $|R\_s|^2$ over the delay-Doppler plane equals the signal energy squared, with the consequence that delay resolution and Doppler resolution cannot both be made arbitrarily narrow for a signal of fixed energy—translates to a corresponding constraint on Zak-domain pulse shapes.

Third, the Zak-domain structure of the Global Navigation Satellite System ranging codes is well-understood in the geodesy literature, though typically not described in Zak-domain terminology. A GNSS pseudorandom noise code of length N chips and chip duration T_c, when considered as a signal s(t) over a full code period, has a Zak transform whose structure determines the code's performance under a combination of delay and Doppler effects; in particular, the GPS L1 C/A code's well-known performance characteristics under Doppler offsets of up to several kilohertz are a consequence of the Zak-domain structure of the Gold-code family from which L1 C/A is constructed [12]. The same mathematical object—the Zak transform of a ranging code—that the geodesy literature uses operationally for GNSS design is the mathematical object that Pillar II treats information-theoretically for beyond-6G waveform design. This is a specific instance of the GNSS–ISAC isomorphism developed in §3.3, specialized to Pillar II.

4.3 The capacity-theoretic status of delay-Doppler channels

The capacity-theoretic status of delay-Doppler channels is widely misunderstood in the applied literature, and the misunderstanding is consequential for the open problems identified below. The clarifying result is that of Lampel, Joudeh, Alvarado and Willems [45], who analyze OTFS in its DZT (discrete Zak transform) formulation and prove that, for the additive-white-Gaussian-noise channel model typically assumed in OTFS analyses, the capacity of the delay-Doppler representation is equal to the capacity of the time-frequency representation of the same underlying channel. This is an equivalence result: the Shannon capacity is a property of the channel, not of the representation in which the channel is described, and different invertible representations must produce the same capacity.

The Lampel equivalence is often read as "OTFS has the same capacity as OFDM," which is true in the setting Lampel et al. analyze but which obscures two important points. The first is that the equivalence depends on the channel model. For a channel described as AWGN plus a deterministic or time-invariant delay-Doppler spreading function, the Lampel equivalence holds. For a channel described as a *stochastic process* on the delay-Doppler plane—in which the locations and amplitudes of the delay-Doppler scatterers are themselves random variables drawn from some ensemble—the capacity question is open, and it is not answered by the Lampel result. This is the distinction between a channel that is time-varying in a deterministic way (Lampel) and a channel that is itself random in a way that the receiver must track across coherence intervals (open). The second is that the equivalence is about scalar capacity, not about the joint sensing-communication Fisher–Shannon metric space $\mathscr{F}_{FS}$ of §3.4. The sensing-side projections of $\mathscr{F}_{FS}$—the Fisher information matrix entries corresponding to delay, Doppler, and complex-amplitude estimation for each scatterer—are representation-dependent in the sense that different representations may make different sensing structure easier or harder to estimate from observations. Lampel does not address the sensing-side projections; his result is purely communication-capacity.

The open question, then, is what a *native* delay-Doppler capacity theorem would look like. Four features are desirable: (i) the channel model is a stochastic process on the delay-Doppler plane, parameterized by a scatterer-number distribution, a delay-Doppler-location distribution, and an amplitude distribution; (ii) the input distribution is optimized over Zak-domain distributions rather than time-frequency-domain distributions, allowing input choices that exploit the sparse structure; (iii) the capacity expression reveals how the sparse-scatterer structure affects the rate-versus-channel-estimation-overhead tradeoff that is intuitively associated with delay-Doppler modulation; and (iv) the capacity expression reduces, in the deterministic-channel limit, to the Lampel-equivalence time-frequency capacity of the same channel realization. No result in the current literature meets all four criteria. The survey treats the derivation of such a theorem as problem O32 in the subcatalog below.

4.4 The ambiguity-function connection and the GNSS signal-design bridge

Sections 4.2 and 4.3 combine to produce the specific bridge that Pillar II offers to the geodetic-estimation-theory toolkit. The bridge has three components.

First, the ambiguity function that the radar literature uses to evaluate a waveform's delay-Doppler resolution is the same object that the Zak transform produces from the waveform's time-domain representation, up to a standard change of variables. The classical radar literature on ambiguity-function design—including the Woodward uncertainty principle, the ambiguity-function volume theorem, and the catalog of waveform families optimized for specific ambiguity-function properties (thumbtack, ridge, Costas)—translates directly into constraints on the Zak-domain pulse shapes that Zak-OTFS uses. The radar literature has been working on this object since 1953; the Zak-OTFS literature encountered it in earnest around 2022; the equivalence has not been exploited to its full extent in the Zak-OTFS literature. This gap is captured by open problems O33 and O34 below.

Second, the GNSS signal-design literature has, since the 1970s, produced a substantial body of results on waveforms optimized jointly for delay resolution and Doppler tolerance—effectively, Zak-domain waveform design, though typically not described in those terms [12], [13]. The design of GPS L1 C/A, L1C, L2C, and L5; the design of Galileo E1, E5a, and E5b; the design of BeiDou B1, B2, and B3; and the joint design of multi-frequency GNSS signals for ionospheric-delay correction (which depends on the ambiguity-function structure across frequency bands) are all, from the Zak-OTFS perspective, worked examples of sophisticated delay-Doppler waveform design. The geodesy-literature record on the performance characteristics of these designs—codeless and semi-codeless tracking, multipath mitigation via correlator spacing, narrow-correlator architectures, strobe correlators, pilot/data combinations for improved tracking thresholds—is a library of empirical results that Pillar II has not mined.

Third, the modulation-level extension of this bridge is direct. A Zak-OTFS pilot-data frame structure in which the pilot symbols are chosen to have GNSS-like autocorrelation properties and the data symbols carry the communication payload is the natural synthesis of GNSS ranging-signal design and Zak-OTFS modulation. Such a waveform is, by construction, simultaneously a communication waveform with the delay-Doppler robustness of Zak-OTFS and a ranging waveform with the code-correlation structure of a GNSS pseudorandom code. The information-theoretic analysis of such a waveform's CRB–Rate region, in the Fisher–Shannon metric space of §3.4, is an open problem (O38 below) that the program's paper P1-6 will address.

4.5 Pillar II open-problem subcatalog

The open problems registered in this module are numbered O32 through O49. Eighteen problems are registered, at the targeted density of 15–20.

O32. Native sparse delay-Doppler capacity theorem. Establish a Shannon-type capacity theorem for channels described as stochastic processes on the delay-Doppler plane, with capacity expression in terms of scatterer-ensemble statistics. Difficulty: Grand Challenge. Prerequisites: information theory for compound channels, Zak-transform analysis, sparse random-process theory. Effort: 24–36 PhD-months. Spiral potential: defines a sub-literature.

O33. Zak-domain Woodward principle. State and prove the Zak-domain analog of Woodward's uncertainty principle, relating Zak-domain pulse-shape support to delay-Doppler resolution. Difficulty: Hard. Prerequisites: time-frequency analysis, Zak transform, uncertainty-principle theory. Effort: 8–12 PhD-months. Spiral potential: 2–3 follow-on papers on pulse-shape-design bounds.

O34. Zak-domain thumbtack pulse classification. Produce a classification of pulse shapes whose Zak transform has thumbtack-like autocorrelation (sharp peak, low sidelobes across the entire delay-Doppler plane), analogous to the classical radar literature on thumbtack waveforms. Difficulty: Hard. Prerequisites: frame theory, pulse-shape optimization, Zak transform. Effort: 8–15 PhD-months. Spiral potential: 3–5 follow-on papers on specific families.

O35. Water-filling in the delay-Doppler domain. Derive the capacity-achieving Zak-domain input distribution for the stochastic delay-Doppler channel of O32, giving an explicit water-filling-type solution. Difficulty: Hard (conditional on O32 being solved). Prerequisites: information theory, convex optimization, Zak-domain analysis. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers on suboptimal structured input distributions.

O36. Delay-Doppler finite-blocklength theory. Develop the Polyanskiy-Poor-Verdú finite-blocklength refinement for delay-Doppler channels, identifying the second-order (dispersion) term. Difficulty: Hard. Prerequisites: finite-blocklength information theory, Zak-domain capacity. Effort: 15–20 PhD-months. Spiral potential: 3 follow-on papers on URLLC-OTFS.

O37. Delay-Doppler MAC capacity region. Characterize the multiple-access capacity region for the delay-Doppler channel with K users communicating on a shared Zak-domain medium. Difficulty: Hard. Prerequisites: MAC capacity, Zak-domain input distributions, sparse-scatterer channels. Effort: 15–20 PhD-months. Spiral potential: 3 follow-on papers on broadcast and interference extensions.

O38. Zak-OTFS CRB–Rate region. Characterize the CRB–Rate region of a Zak-OTFS ISAC system with embedded-pilot structure, recovering the GNSS ranging-CRB in the pilot-only limit. This is the program's P1-6 paper. Difficulty: Hard. Prerequisites: Zak-OTFS, CRB-rate region (§3), GNSS code design. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers on specific ISAC-OTFS settings.

O39. Zak-domain ambiguity-function synthesis. Translate the classical radar-literature ambiguity-function catalog (Barker, Costas, P3, P4, Frank, Zadoff-Chu) into the Zak-OTFS pulse-shape context, identifying which classical ambiguity-function constructions have Zak-OTFS analogs. Difficulty: Accessible. Prerequisites: radar ambiguity function theory, Zak transform, pulse-shape design. Effort: 6–12 PhD-months. Spiral potential: 3–5 follow-on papers on individual constructions.

O40. Delay-Doppler capacity under fractional Doppler. Extend the Zak-OTFS capacity analysis to settings in which the Doppler shift per scatterer is not an integer multiple of the Zak-grid Doppler spacing, producing capacity expressions that account for fractional-Doppler-induced inter-symbol interference. Difficulty: Hard. Prerequisites: Zak transform, interpolation theory, information-theoretic capacity. Effort: 12–18 PhD-months. Spiral potential: 2–3 follow-on papers.

O41. Learned Zak-domain equalization. Characterize the performance of neural-network equalizers operating in the Zak domain, as a function of the number of scatterers, the scatterer-location distribution, and the neural-network parameter count. Difficulty: Hard. Prerequisites: neural equalization, PAC learning, Zak-domain channel models. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers at the Pillar II / Pillar IV interface.

O42. Crystallization condition under hardware constraints. Characterize how realistic hardware constraints (finite-word-length digital-to-analog converters, pulse-shape implementation tolerances, phase noise) modify the Zak-OTFS crystallization condition. Difficulty: Accessible–Hard. Prerequisites: Zak-OTFS, hardware-impaired signal processing. Effort: 8–12 PhD-months. Spiral potential: 3–4 follow-on papers.

O43. Delay-Doppler capacity under LEO channel models. Specialize the delay-Doppler capacity analysis to LEO-satellite channel models, where Doppler magnitudes are large and Doppler rates are substantial within a single frame. Difficulty: Hard. Prerequisites: Zak-domain capacity, LEO channel modeling. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers at the Pillar II / Pillar III interface.

O44. Delay-Doppler secrecy capacity. Characterize the secrecy capacity of a delay-Doppler channel with an eavesdropper whose channel differs from the legitimate receiver's in the sparse-scatterer structure. Difficulty: Hard. Prerequisites: physical-layer secrecy, Zak-domain capacity. Effort: 12–18 PhD-months. Spiral potential: 2–3 follow-on papers.

O45. Delay-Doppler capacity with mmWave beamforming. Combine the Zak-OTFS framework with millimeter-wave beamforming, producing capacity expressions that exploit both the sparse angular support at mmWave and the sparse delay-Doppler support. Difficulty: Hard. Prerequisites: Zak-domain capacity, mmWave channel modeling, hybrid beamforming. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers.

O46. Pulse-shape-Zak joint design for mixed data/sensing frames. Derive jointly optimal Zak-domain pulse shapes and delay-Doppler symbol distributions for mixed data/sensing frames, where part of the frame carries communication data and part performs integrated sensing.

Difficulty: Hard. Prerequisites: Zak-OTFS, joint design of pilots and data, Fisher–Shannon metric space (§3.4). Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O47. Capacity-gap characterization between OTFS and Zak-OTFS. For specific channel models, characterize the capacity gap between the original Hadani OTFS formulation (with imperfect pulse-shape compatibility) and the Zak-OTFS formulation. Difficulty: Accessible. Prerequisites: OTFS, Zak-OTFS, capacity analysis. Effort: 6–10 PhD-months. Spiral potential: 2 follow-on papers clarifying the practical implications.

O48. Delay-Doppler capacity with imperfect CSI. Extend the Zak-domain capacity analysis to settings where channel-state information is imperfect at the transmitter, receiver, or both, and characterize the capacity penalty as a function of the CSI error statistics. Difficulty: Hard. Prerequisites: imperfect-CSI capacity, Zak-domain capacity. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O49. Delay-Doppler capacity in the presence of a non-cooperative sensor. Characterize the capacity of a Zak-OTFS communication link in the presence of an uncooperative sensing system that is trying to extract information from the waveform, with implications for covert and low-probability-of-intercept communications. Difficulty: Hard. Prerequisites: covert communication, Zak-domain signal analysis. Effort: 12–18 PhD-months. Spiral potential: 2–3 follow-on papers.

These eighteen problems span the capacity-theoretic core of Pillar II (O32, O35, O36, O37), the pulse-shape and waveform-design bridge to radar and GNSS (O33, O34, O39, O42, O46), the ISAC interface (O38, O46), the AI-native interface (O41), the Pillar III interface (O43), and the realistic-system directions (O40, O42, O44, O45, O47, O48, O49).

5. Pillar III — SAGIN Positioning-Communication Integration

The third pillar treats Space-Air-Ground Integrated Networks—the heterogeneous multi-layer networks in which geostationary, medium-Earth-orbit, and low-Earth-orbit satellites are combined with high-altitude platforms, unmanned aerial vehicles, and terrestrial base stations to provide continuous coverage with spatial diversity across altitudes—as a wireless pillar whose information-theoretic foundations are largely unbuilt. The pillar's central diagnostic observation is that the SAGIN literature, which is extensive and active, has concentrated almost exclusively on architectural and protocol-level questions (routing, handover, resource orchestration, slicing, network-function placement), and has produced correspondingly few results at the level at which Pillars I, II, IV and V have been working: capacity regions, Cramér–Rao bounds on joint positioning-communication performance, Fisher-information decompositions across layers. The gap is a specific instance of a discipline-level mismatch: the multi-constellation, multi-layer, geometry-dependent estimation-theoretic machinery that Pillar III needs exists in the geodesy literature under the heading of "multi-constellation GNSS" and has been developed operationally over two decades of GPS–GLONASS–Galileo–BeiDou interoperability work. The present section reviews the SAGIN literature explicitly as architectural, reviews the LEO-PNT literature as the closest existing analog to a SAGIN positioning-theoretic tradition, and maps the estimation-theoretic machinery into the open problems that define Pillar III. Figure 5.1 provides the architectural reference picture for the section: a multi-layer SAGIN with the Fisher-information decomposition across layers annotated explicitly.

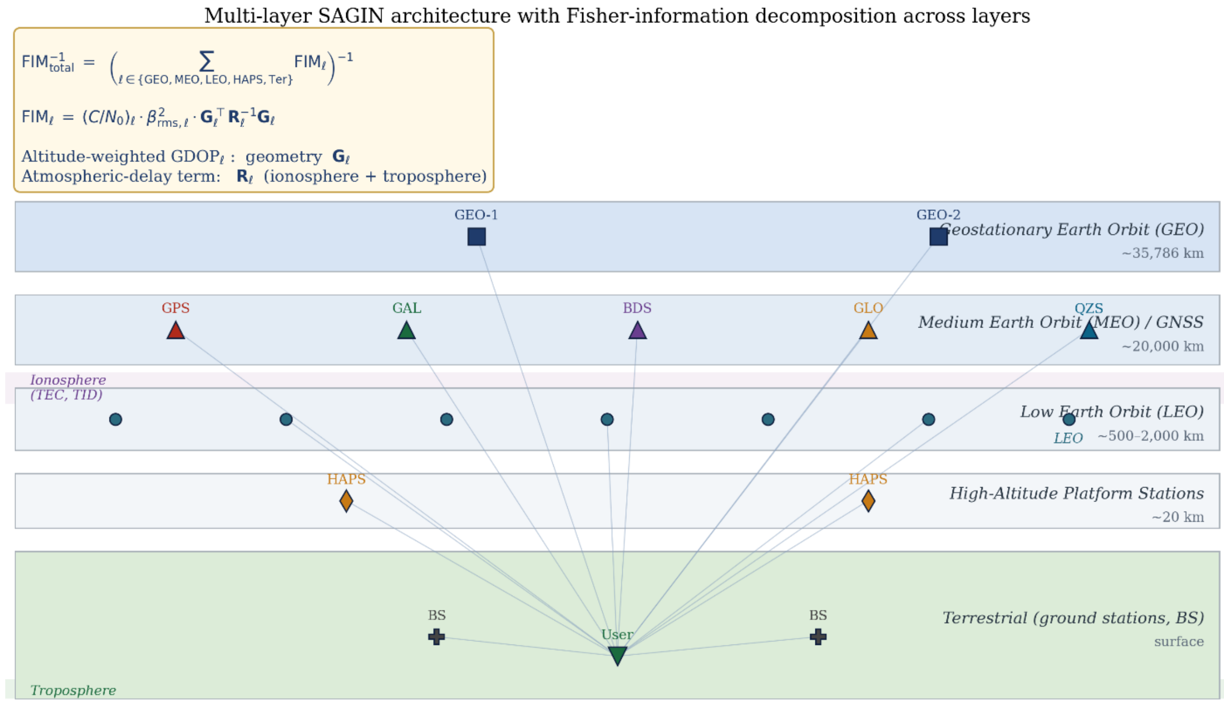

Figure 5.1. Multi-layer SAGIN architecture with Fisher-information decomposition across layers. Each altitude layer (GEO, MEO/GNSS, LEO, HAPS, terrestrial) contributes an additive term FIM_ℓ to the joint Fisher information matrix on user state, with per-layer geometry G_ℓ, signal-strength factor (C/N_0)_ℓ, root-mean-square bandwidth β_{rms,ℓ}, and atmospheric-delay covariance R_ℓ. The total positioning covariance $\mathrm{FIM\_total}^{-1}$ is the inverse of the sum across layers; the altitude-weighted GDOP formalism (problem O51) is the SAGIN-specific generalization of the classical single-layer GNSS DOP. The ionosphere band (between GNSS/MEO and LEO) and troposphere band (at the surface) are explicit measurement-model elements that couple positioning accuracy, communication-decoding performance, and carrier-phase tracking stability across all layers.

5.1 The SAGIN literature as architectural

The SAGIN literature in its modern form begins with the surveys of Liu, Shi and collaborators [34], who framed SAGIN as the unification of satellite networks, aeronautical networks, and terrestrial cellular networks, and who identified the principal engineering challenges—mobility management, routing across heterogeneous layers, resource management across links with disparate propagation characteristics, security across network segments owned by different operators. The Liu-Shi line and its contemporaries established the vocabulary and the problem taxonomy for SAGIN and set the pattern that subsequent surveys have followed: a layered architectural framing, a catalog of enablers and challenges, and a set of cross-layer optimization problems stated at the level of integer programming or convex optimization rather than at the level of capacity theory.

The service-oriented 6G SAGIN survey of Cheng, He, Yin and collaborators [35] extends the framework to explicitly 6G timescales, identifying service classes (enhanced mobile broadband, ultra-reliable low-latency, massive machine-type, and their 6G successors) and matching them to

SAGIN capabilities. The AI-SAGIN interplay survey from KAUST [59] consolidates the machine-learning-for-SAGIN literature—resource-allocation policies learned from operational data, routing policies learned via reinforcement learning, handover policies learned across traffic regimes. The joint-communication-and-computing survey by Chen, Guo, Meng, Han, Li and Quek [36] extends SAGIN to include computational resources distributed across the layers, producing a joint communication-computing-networking optimization framework. The in-flight-6G-slice-orchestration work of Toka et al. [60] addresses the specific problem of supporting 6G network slices across the ground-to-aircraft segment, with explicit resource-dimensioning analyses. The Zhang et al. article [61] and the broader space-air-ground survey [62] provide architectural overviews; the Kodheli et al. survey [63] provides the detailed non-terrestrial-network background.

The common property of these surveys is that positioning does not appear as a performance axis of the SAGIN capacity optimization. It appears as a use case the SAGIN must support (e.g., "positioning applications require 10 ms latency to the edge" appears in service-class specifications), but it does not appear as a quantity to be jointly optimized against communication rate in the way that Pillars I and II treat it. The capacity questions the surveys address are communication-capacity questions across the heterogeneous topology; the positioning questions are treated, when they are treated at all, as separate and subsequent.

A second common property is that the mathematical machinery used is that of optimization and graph theory rather than that of estimation theory. A SAGIN resource-allocation problem is typically formulated as a mixed-integer linear program or a convex optimization with capacity constraints, routing constraints, and delay constraints; it is solved by general-purpose or specialized convex-optimization algorithms. The Cramér–Rao bound, the Fisher information matrix, the Kalman filter, and the dilution-of-precision formalism—the principal tools of Pillars I, II, IV and V and of the geodesy literature that Pillar III's opportunity rests on—do not appear in the SAGIN surveys' index of methods. This is a choice of framing, not a mathematical limitation of SAGIN; it reflects which research community has produced the surveys.

5.2 The LEO-PNT literature as the closest analog

The closest analog to an estimation-theoretic SAGIN literature is the LEO-PNT (low-Earth-orbit positioning, navigation and timing) literature, which has grown substantially over the past decade as the proliferation of LEO communication constellations—Starlink, OneWeb, Kuiper, Lightspeed, the Chinese SatNet megaconstellation, the Iridium Next modernized system—has made LEO signals available as signals-of-opportunity for positioning. The Khalife, Kassas, Morales and Ragen line of work [64], which has produced a series of experimental and analytical results on the positioning performance achievable using commercial LEO downlink signals without the cooperation of the constellation operator, is the principal research program in the area. Reid, Walter and collaborators [65], [66] have produced analytical bounds on LEO-PNT accuracy as a function of constellation geometry, signal bandwidth, and user receiver architecture. The European Space Agency's Moonlight Lunar Communications and Navigation System initiative [67] represents the first institutional commitment to extend PNT infrastructure beyond Earth orbit, with explicit LEO-analog architectural choices.

The LEO-PNT literature is estimation-theoretic in its style and in its methods. It uses the Cramér–Rao bound to characterize ranging accuracy; it uses dilution-of-precision to characterize the geometric factor; it uses the Kalman filter and extended Kalman filter to characterize position tracking; and it treats atmospheric delays (ionospheric and tropospheric) as explicit measurement-model terms. In other words, the LEO-PNT literature uses the geodesy toolkit directly, because its practitioners come largely from the GNSS community and are working on what is, from their vantage, a natural extension of terrestrial GNSS into the LEO-downlink regime.

The LEO-PNT literature is, however, narrowly scoped. It treats the LEO signal as a one-way downlink from a satellite whose position is known to a receiver whose position is unknown; the communication content of the signal (when the signal has communication content) is treated as a nuisance variable rather than as a performance axis. The information-theoretic question—what is the joint capacity-positioning region of a LEO signal used simultaneously for communication and PNT?—is not standard in the LEO-PNT literature. Neither is the multi-layer question: what is the joint capacity-positioning region when LEO signals are combined with MEO (GNSS) signals, with GEO communication signals, with HAPS signals, and with terrestrial 6G signals, as a reader of a SAGIN architectural survey would naturally ask?

Pillar III's opportunity, then, is that the estimation-theoretic toolkit exists in the LEO-PNT community, but has not been extended to the capacity-positioning joint setting; and the capacity-and-architecture toolkit exists in the SAGIN community, but has not been extended to the estimation-theoretic positioning setting. The pillar's research program, as catalogued in §5.5, is to perform the extension from both sides.

5.3 Multi-constellation estimation and dilution of precision across layers

The mathematical machinery that Pillar III's opportunity rests on is the multi-constellation estimation theory of GNSS. Multi-constellation GNSS—the combination of GPS, GLONASS, Galileo, BeiDou, QZSS, IRNSS (now NavIC) and various regional augmentations into a single receiver solution—has been operational for approximately two decades, with the performance, interoperability and systematic-error characteristics understood in detail [12], [13], [14]. The mathematical apparatus has three relevant components.

The first component is the extension of the Cramér–Rao bound to multi-constellation geometries. For a receiver simultaneously tracking N satellites from M constellations, with N_m satellites from constellation m (so $N = \Sigma_m N_m$), the inverse Fisher information matrix on the position estimate is the sum of contributions from each satellite, weighted by the inverse of the per-satellite measurement variance. The per-satellite measurement variance depends on the per-satellite ranging CRB, on the atmospheric-delay measurement accuracy for the satellite's line-of-sight path, and on the systematic inter-constellation timing offsets that must be estimated jointly with the position. The structure is a direct generalization of the single-constellation GDOP, with additional dimensions in the estimation state-vector for inter-constellation offsets and inter-frequency biases.

The second component is the weighting of multi-constellation contributions for robustness. When one constellation is degraded—ionospheric scintillation affecting high-latitude GPS signals, jamming or spoofing affecting L-band, deliberate signal-structure changes by constellation operators—the receiver's position estimate should be robust against the degradation. The geodesy literature has developed weighting schemes (inverse-variance weighting, M-estimators, robust Kalman filters) that produce position estimates that remain accurate under partial degradation. The author's own work on GPS detection of the Minuteman-II launch from Vandenberg Air Force Base on 8 July 2000 [15] is an operational example of estimation under a sparse and degraded observation set: the TEC perturbation released by the launch—pulses exceeding 6×10^{14} electrons per square metre—was extracted from seven stations of a commercial GPS network by dual-frequency geometry-free combination and high-pass filtering, and the launch site was subsequently recovered to within 35 m by weighted least squares over the resulting subionospheric-point geometry. The same weighted-estimation machinery generalizes directly to the multi-constellation case, where the weights absorb inter-constellation timing offsets alongside per-satellite measurement variance.

The third component is the atmospheric-delay-correction formalism. The ionospheric delay on a GNSS signal is, to first order in the carrier frequency, proportional to the total electron content (TEC) along the line-of-sight path and inversely proportional to the square of the carrier frequency. Multi-frequency GNSS (dual-frequency at minimum, triple-frequency increasingly) uses this frequency dependence to solve for the ionospheric delay as part of the measurement model rather than treating it as residual noise. The tropospheric delay has a standard-atmosphere component that is removed by empirical models (Saastamoinen, Hopfield) and a residual component that is estimated jointly with the position in high-accuracy receivers. The formalism is compact and well-understood [12].

The relevance of this formalism to SAGIN is structural. A LEO-MEO-GEO SAGIN receiver combining signals from multiple altitudes faces exactly the mathematical situation of a multi-constellation GNSS receiver: signals from geometries with different geometric-dilution properties, with altitude-dependent atmospheric delays, and with inter-operator timing offsets that must be estimated jointly with user position and velocity. The multi-constellation GNSS machinery is the mathematically correct tool for the problem; it has been deployed operationally for two decades in one context; and its extension to SAGIN altitudes and to joint positioning-communication optimization is an open problem set (O50–O57 in the subcatalog).

5.4 Atmospheric delay as an information-theoretic constraint

The second geodesy-toolkit contribution Pillar III brings is the formal treatment of atmospheric delay as an information-theoretic constraint on SAGIN performance. The point is conceptually simple but has not been made in the SAGIN survey literature: for any SAGIN signal that propagates through the ionosphere or lower atmosphere, the atmospheric delay is a measurement-model term that couples to the positioning accuracy, the communication decoding performance, and the carrier-phase tracking stability in ways that the SAGIN surveys have not analyzed. The operational geodesy literature has characterized these couplings thoroughly; the wireless-communications SAGIN literature has not.

The author's work on detecting the traveling ionospheric disturbance excited by the Minuteman-II launch from Vandenberg Air Force Base on 8 July 2000 [68], using the dense Southern California Integrated GPS Network (135 permanent stations spanning 591 km) as a distributed TEC sensor, is a concrete example. The ionospheric TID signature—a spatially-coherent TEC perturbation appearing across the network on average fifteen minutes after launch, propagating within the ionospheric layer at approximately 1230 m/s and coherent for some forty-five minutes—was detectable at the sub-TECU level against the ambient ionospheric variability only because the GPS-network TEC estimation machinery provides simultaneous position, carrier-phase, and ionospheric-delay estimates at each station with uncertainties characterized by the Fisher-information structure described in §5.3. The information-theoretic content of the observation is that a SAGIN whose LEO or GEO downlink signals were subject to the same ionospheric TID would experience a correlated phase perturbation at the sub-nanosecond level across widely-separated ground stations—a correlation structure that is simultaneously a positioning-accuracy nuisance (for applications requiring sub-decimeter positioning) and a potentially-exploitable signal-structure feature (for applications seeking to characterize ionospheric conditions from the SAGIN downlink itself). That the attribution survived at all is itself an estimation-theoretic result: the apparent westward propagation, toward rather than away from the launch site, was reconciled only by admitting multi-altitude ionospheric pierce points, and competing sources—solar flares and seismic activity—were excluded on independent geomagnetic and seismographic evidence. The SAGIN capacity-positioning region, properly formulated, would account for both aspects.

The LEO-PNT literature has begun treating atmospheric delay in the LEO-downlink context; the SAGIN communications literature has not. The formal characterization of atmospheric delay as an entry in the Fisher information matrix for joint positioning-communication is an open problem (O59 below). The characterization of how ionospheric scintillation—well-studied in the GNSS literature—affects SAGIN joint-decoding performance across LEO and MEO signals is another open problem (O60). The ICL-for-navigation work of the author's group [17] provides a machine-learning angle on these questions: an in-context-learning framework for GNSS/navigation that adapts to changing ionospheric conditions without retraining is the natural AI-native substrate for a SAGIN receiver operating across an atmosphere whose scintillation statistics change rapidly.

5.5 Pillar III open-problem subcatalog

The open problems registered in this module are numbered O50 through O67. Eighteen problems are registered.

O50. Multi-constellation SAGIN CRB–Rate region. Establish the joint capacity-positioning region for a multi-layer SAGIN receiver combining signals from LEO, MEO, GEO, HAPS and terrestrial sources, with explicit Fisher-information decomposition across layers. Difficulty: Grand Challenge. Prerequisites: multi-constellation estimation, CRB-rate region (§3), SAGIN architecture. Effort: 18–30 PhD-months. Spiral potential: defines the Pillar III central sub-literature; the program's P1-5 paper is the first entry.

O51. Altitude-weighted GDOP formalism. Derive a SAGIN-specific altitude-weighted GDOP formalism that incorporates the per-altitude atmospheric-delay uncertainty, per-altitude ranging-signal-structure, and per-altitude Doppler-rate characteristics into a single performance scalar. Difficulty: Hard. Prerequisites: GNSS GDOP, Fisher-information decomposition, SAGIN architecture. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers on specific altitude-combination scenarios.

O52. LEO-MEO capacity-positioning integration. Characterize the joint capacity-positioning region for a receiver combining LEO communication-signals-of-opportunity with MEO GNSS signals, with Fisher-information decomposition showing the LEO contribution to the positioning CRB and the MEO contribution to the communication rate. Difficulty: Hard. Prerequisites: LEO-PNT, GNSS, CRB-rate region. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers on specific constellation combinations.

O53. HAPS-terrestrial handover capacity-positioning continuity. Characterize the joint capacity-positioning continuity as a SAGIN user transitions between high-altitude-platform coverage and terrestrial-base-station coverage, with positioning-CRB and communication-rate continuity conditions. Difficulty: Hard. Prerequisites: mobility management, positioning continuity, capacity analysis. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O54. Inter-constellation timing-offset capacity implications. Quantify the capacity and positioning implications of inter-constellation timing offsets in a multi-layer SAGIN receiver, characterizing the tradeoff between estimating these offsets online and taking them as fixed calibration parameters. Difficulty: Accessible–Hard. Prerequisites: multi-constellation timing, Kalman filtering, capacity analysis. Effort: 8–12 PhD-months. Spiral potential: 2–3 follow-on papers.

O55. SAGIN Kalman-filter steady-state covariance. Derive the steady-state covariance of a multi-layer SAGIN Kalman filter tracking joint position, velocity, and communication-channel-state in a recurring geometry. Difficulty: Hard. Prerequisites: Kalman-filter theory, SAGIN geometry, multi-layer signal combination. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers on specific maneuver scenarios.

O56. SAGIN robustness under constellation degradation. Characterize the capacity-positioning region degradation as layers of the SAGIN fail or are degraded (LEO outage, regional GNSS denial, terrestrial jamming), producing a continuous degradation curve rather than a discrete failure mode. Difficulty: Hard. Prerequisites: robust estimation, SAGIN capacity-positioning. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers.

O57. Signals-of-opportunity information-theoretic characterization. Formalize the information-theoretic content of signals-of-opportunity for positioning—characterizing the extractable Fisher information as a function of the non-cooperative-signal structure, the receiver's prior knowledge, and the observation duration. Difficulty: Hard. Prerequisites: non-cooperative signal estimation, Fisher-information theory. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O58. SAGIN capacity under Doppler-rate effects. Extend the SAGIN capacity analysis to explicitly account for Doppler rate (not just Doppler shift) within a coherent frame, which is substantial for LEO downlinks and which interacts with Pillar II's delay-Doppler analysis (cross-reference O43). Difficulty: Hard. Prerequisites: Doppler-rate modeling, Zak-domain capacity, LEO links. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O59. Atmospheric delay as Fisher-information-matrix entry. Derive the contribution of atmospheric-delay uncertainty (ionospheric TEC residual, tropospheric wet-delay residual) to the Fisher information matrix for joint SAGIN positioning-communication, explicitly showing how multi-frequency signal combinations reduce the atmospheric-uncertainty contribution. Difficulty: Hard. Prerequisites: atmospheric-delay modeling, Fisher-information theory, multi-frequency GNSS. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers.

O60. Ionospheric scintillation effects on SAGIN joint decoding. Characterize the degradation of SAGIN joint communication-decoding and positioning performance under ionospheric scintillation, with scintillation statistics parameterized by the standard ionospheric-scintillation indices. Difficulty: Hard. Prerequisites: ionospheric scintillation, joint decoding theory. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O61. Multi-frequency SAGIN capacity-positioning. Characterize the capacity-positioning region when SAGIN signals span multiple frequency bands (L-band, C-band, Ku-band, Ka-band, V-band), with explicit frequency-dependent atmospheric-delay, rain-attenuation, and antenna-pattern effects. Difficulty: Hard. Prerequisites: multi-frequency signal combination, SAGIN capacity, atmospheric propagation. Effort: 12–18 PhD-months. Spiral potential: 4–6 follow-on papers on specific band combinations.

O62. ESA LCNS / Moonlight-style beyond-Earth SAGIN theory. Extend the SAGIN capacity-positioning framework to include lunar-orbit and cislunar navigation constellations, characterizing how the SAGIN framework applies to navigation architectures beyond Earth orbit. Difficulty: Hard. Prerequisites: SAGIN capacity-positioning, space navigation. Effort: 18–24 PhD-months. Spiral potential: 3–5 follow-on papers on deep-space SAGIN.

O63. SAGIN privacy constraints on positioning. Characterize the joint capacity-positioning region under privacy constraints that limit the sensing information leaking to the network about user positions (differential privacy applied to positioning estimates). Difficulty: Hard. Prerequisites: differential privacy, positioning CRB, SAGIN capacity. Effort: 12–18 PhD-months. Spiral potential: 2–3 follow-on papers.

O64. SAGIN capacity under spectrum-sharing constraints. Characterize the capacity-positioning region when the SAGIN operates in spectrum shared with other services (primary users, other SAGIN operators), with explicit interference constraints. Difficulty: Hard. Prerequisites: cognitive-radio capacity, SAGIN architecture. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O65. Energy-efficient SAGIN capacity-positioning. Extend the SAGIN capacity-positioning region to include explicit energy-consumption axes across the heterogeneous satellite, platform, and terrestrial transmitters. Difficulty: Hard. Prerequisites: energy-efficient capacity, SAGIN. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O66. AI-native SAGIN receiver sample complexity. Characterize the performance of neural-network SAGIN receivers (specifically in-context-learning architectures for multi-constellation positioning) as a function of the receiver parameter count, training-data diversity across atmospheric conditions, and geometry-distribution support. Difficulty: Hard. Prerequisites: in-context-learning theory, SAGIN capacity-positioning, Pillar IV theory. Effort: 15–20 PhD-months. Spiral potential: 3–4 follow-on papers at the Pillar III / Pillar IV interface.

O67. SAGIN capacity under GNSS denial. Characterize the capacity-positioning region achievable by a SAGIN receiver when GNSS is wholly denied (jammed or spoofed to unusability), and must rely on LEO signals-of-opportunity, terrestrial ranging signals, and inertial integration. Difficulty: Hard. Prerequisites: signals-of-opportunity positioning, alternative PNT, SAGIN. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers on specific alternative-PNT architectures.

Eighteen problems span the multi-layer capacity-positioning core (O50–O52, O53, O55, O56), the signals-of-opportunity question (O57, O67), the atmospheric and ionospheric questions (O59, O60, O61), the beyond-Earth extension (O62), the AI-native interface (O66, continuing O41's trajectory), and the realistic-constraint directions (O54, O58, O63–O65). The three cross-pillar interfaces—Pillar II (O58, O43 cross-ref), Pillar IV (O66), and the survey's Section 9 cross-pillar section (O62 previewing the quantum-PNT convergence)—are explicitly marked.

6. Pillar IV — AI-Native Physical Layer Theory

The fourth pillar treats what the standards literature has come to call the AI-native physical layer: the systematic incorporation of learned components—autoencoder-based modulation-and-coding chains, neural receivers for demodulation and channel estimation, learned beam-management policies, foundation-model-based channel-state compression and prediction—into the physical layer of a wireless system as first-class components rather than as optional add-ons. The engineering momentum behind AI-native PHY is substantial: 3GPP has moved AI/ML-for-PHY from study items in Release 18 toward normative specifications in Release 19 and 20, every major wireless vendor has demonstrated neural-receiver prototypes at trade shows, and the large-scale industrial software platforms (NVIDIA Sionna, DeepSig OmniPHY, Samsung AI-PHY, Rohde & Schwarz neural-receiver tools, MathWorks 6G AI blocks) are now standard elements of the physical-layer development environment. The theoretical situation, however, is markedly less developed. Pillar IV's central diagnostic is that the engineering momentum has outrun the foundational theory: a body of specifications and prototypes is accumulating whose performance claims are empirical rather than provable, whose safety and robustness properties are largely uncharacterized, and whose design decisions are consequently difficult to port across deployment scenarios, to extend to new use cases, or to defend against adversarial manipulation. The pillar's research program is to construct the theoretical substrate that the engineering prototypes currently lack, with emphasis on capacity-theoretic bounds for learned codebooks, on sample-complexity bounds for learned receivers, and on provable-guarantee analogs for AI-native estimation that play the role the Kalman covariance matrix plays for classical estimation.

6.1 The autoencoder-as-communication foundational line

The conceptual foundation of AI-native physical-layer theory is the autoencoder-as-communication framework introduced by O'Shea and Hoydis [69] and developed by Dörner, Cammerer, Hoydis and ten Brink [70]. The framework treats an end-to-end communication chain as a single differentiable computational graph: a transmitter neural network that maps information bits to channel-input symbols, a channel model that applies noise and distortion, and a receiver neural network that maps channel outputs back to decoded bits. The system is trained jointly by back-propagating the decoding-error gradient through the channel model to the transmitter network, producing constellation-and-coding pairs that are optimized for the channel in a single end-to-end sense. The framework is significant in three respects.

First, it replaces the design-by-information-theory-followed-by-signal-processing pipeline (Shannon-capacity analysis, then code design, then modulation-and-coding-scheme selection, then decoder implementation) with a single differentiable optimization. The pipeline is reduced to a gradient-descent problem on a parameterized family of encoder-decoder pairs, with the channel model as a fixed but differentiable intermediate. For channels that deviate from the canonical AWGN or standard-fading models (hardware-impaired channels, nonlinear-amplifier channels, channels with non-Gaussian noise distributions), the framework often produces encoder-decoder pairs that outperform hand-designed pipelines.

Second, it exposes a fundamental tension that Pillar IV's theoretical program centers on: the learned encoder-decoder pair is, by construction, optimized for a specific channel model (the one used during training). When the deployment channel deviates from the training channel, the learned system's performance degrades in ways that the information-theoretic capacity analysis of the deployment channel does not bound. The gap between training-channel performance and deployment-channel performance is the quantity that classical statistical-learning theory calls “generalization gap,” and the bounds that statistical-learning theory produces for generalization gaps are in principle applicable to the AI-native PHY setting. In practice the bounds are rarely tight and almost never translated into standards-usable performance guarantees.

Third, the framework naturally extends to semantic communication, where the encoder-decoder pair is optimized not to recover the transmitted bits exactly but to preserve a task-defined semantic content of the source, with the channel rate measured in task-relevant rather than bit-exact terms. The semantic-communication literature [71] has argued, plausibly, that semantic communication can exceed the classical Shannon rate for specific task-metric choices; the information-theoretic status of those claims is incomplete.

6.2 Industry momentum and standards context

The industrial and standards momentum for AI-native PHY is, as of April 2026, considerable and moving faster than the foundational theory. 3GPP Release 18, which completed its specification freeze in 2024, introduced AI/ML for the physical layer through a series of study items covering channel-state-information feedback compression, beam-management policy learning, and positioning model combination [21]. Release 19, in progress at the time of writing, extends these study items to the normative-specification stage for specific use cases. Release 20, scheduled for specification freeze in 2027 and widely identified as the first "6G" release, is expected to contain the first normative AI-native PHY components; Release 21, scheduled for approximately 2028, is expected to contain the first full AI-native air interface.

The industrial support for this trajectory is coordinated and well-resourced. NVIDIA's Sionna platform [72] provides differentiable implementations of 5G New Radio physical-layer chains, including neural-receiver variants, as open research software. DeepSig's OmniPHY [73] offers commercial neural-receiver implementations with vendor-demonstrated performance gains in specific channel regimes. Rohde & Schwarz, Samsung Research, Keysight, and MathWorks have announced or demonstrated AI-PHY tools during 2025 and early 2026 [37], [74]. The ETH Zurich and NVIDIA collaboration on Deep-Unfolded Iterative Detector-Decoder architectures [75] represents the academic side of the same trajectory, producing neural-receiver designs whose architecture is derived from unrolling classical iterative detection-decoding algorithms.

The recent synthesis literature includes the Science China Information Sciences overview by Cui et al. [38], which places AI-native PHY in the context of the full 6G AI stack, and the Frontiers comprehensive review of AI-native 6G integrating semantic communications, reconfigurable intelligent surfaces, and edge intelligence [39]. These surveys are valuable as landscape maps but are explicitly not foundational-theory surveys. They identify AI-native PHY as an engineering trajectory, catalog the enablers, and note the absence of foundational bounds as a future-work item rather than treating that absence as the central theoretical problem.

The imbalance between industrial activity and foundational theory is, in the view of this survey, the most consequential gap among the five pillars, because the standards decisions of Release 19, 20 and 21 will bake engineering-empirical choices into normative specifications on a timescale that does not wait for foundational theory to catch up. A foundational theorem for AI-native PHY published in 2028 will inform textbook treatments and future generations; it will not inform Release 21. A theorem published in 2026 or 2027 will inform Release 21 and influence Release 22. Pillar IV's timing argument, stated in §1.3, is sharpest here.

6.3 The theoretical voids

Four theoretical questions stand out as the foundational problems of Pillar IV. The four correspond to the four places in a classical Shannon-plus-estimation analysis at which substituting a learned component for a classical one produces a question to which the classical theory does not answer.

Neural-codebook capacity. The first question is capacity under learned-encoder constraints. Shannon capacity assumes the encoder ranges over all codebooks of a given rate; a learned encoder ranges only over the codebooks expressible by the parameterized function class (typically a neural network of specified architecture and parameter count). The capacity-under-learned-encoder constraint is, for a general channel, bounded above by the Shannon capacity of the channel; the gap between the two is the "neural-codebook capacity gap" and is a function of the architecture and parameter count. Partial results exist for specific architectures (autoencoder with bounded-weight fully-connected layers, specific convolutional families), but no general theorem relates the gap to the architecture complexity in a way that would allow a standards body to specify neural-encoder complexity requirements sufficient for a given capacity target. Problem O68 in the subcatalog frames this question formally.

PAC-learning bounds for neural receivers. The second question is sample complexity. A learned receiver is a function class parameterized by a training data set of channel-input-output pairs; statistical-learning theory (Vapnik-Chervonenkis, Rademacher-complexity, PAC-Bayes) provides bounds on the number of training samples required for the learned receiver to achieve within ε of the function-class-optimal performance with probability at least $1 - \delta$. The bounds, applied to practical neural-receiver architectures, typically produce sample-complexity numbers that exceed realistic training-dataset sizes by orders of magnitude. The question is which of the mathematical relaxations available in the literature (Rademacher complexity over the achievability set, algorithmic-stability bounds, flat-minima bounds) produces sample-complexity estimates that are simultaneously tight enough to be informative and realistic enough to be achievable, for the specific neural-receiver architectures that standards bodies are considering. The author's work on in-context-learning characterization [76] provides a partial angle on this question: in-context learning is a setting in which the sample-complexity question takes the form of “how many context examples are needed for the model to reach within ε of the Bayes-optimal posterior on the task,” and the theoretical characterization produced in that work is directly relevant to in-context-learning-based receiver architectures. Problem O69 frames the PAC-learning question as an ISAC-theoretic specialization; O19 (from Module 3) frames its Pillar-I interface.

Convergence guarantees for end-to-end autoencoder training. The third question is training convergence. An autoencoder-communication chain is trained by stochastic gradient descent on a loss function whose landscape is non-convex in the encoder and decoder parameters jointly. Whether the training converges to a useful local minimum, and whether the convergence is robust to the initialization and to the training-data ordering, are empirical questions at present. Theoretical results for gradient-descent convergence on specific neural-network families (neural tangent kernel, infinite-width limits, mean-field limits) have developed substantially since 2018; their adaptation to the communication-chain setting is incomplete. Problem O70 frames this question.

Provable-safety analogs for AI-native estimation. The fourth question—the one most central to the pillar's distinctive contribution—is the provable-safety analog for AI-native estimation. A classical Kalman filter, properly matched to the measurement model, produces covariance estimates whose correctness is a theorem: the estimated covariance is the actual covariance of the state-estimate error, up to the measurement-model specification. A neural estimator produces an implicit posterior whose correctness is, in general, not a theorem. For applications in which the estimator's output is consumed by safety-critical downstream logic (autonomous driving, aircraft control, medical-device actuation, aviation navigation), the absence of a provable-safety analog is not merely a theoretical inconvenience; it is the obstacle to fielding AI-native estimators in safety-critical settings. The author's critical analysis of AI-safety red-teaming practices [77] argues that the current industry focus on empirical red-teaming as the primary safety methodology is insufficient for applications where correctness must be provable rather than empirical. The Pillar IV research program treats the construction of provable-safety analogs as the highest-priority AI-native-PHY problem for safety-critical applications. Problem O71 frames it as a formal research question.

6.4 The in-context-learning line and foundation models for navigation

The author's research line on in-context learning and on foundation models for navigation and geospatial applications supplies concrete substrate for several of the Pillar IV open problems. This subsection describes the line briefly and places its entries against the open problems where they are most directly relevant.

In-context learning and its characterization. In-context learning, in the transformer-model sense, is the phenomenon whereby a large pre-trained model produces task-specific outputs at inference time based on a small number of task examples provided in the input context, without updating model weights. The theoretical status of the phenomenon has been substantially clarified in recent years: in-context learning is, in a precise sense, implementing implicit Bayesian inference over a task-distribution prior that was induced by the pre-training corpus, with the context examples playing the role of observations that update the posterior. The author's characterization of in-context learning [76] provides a specific framework for analyzing when and how in-context learning succeeds for tasks outside the pre-training distribution—a setting that is central to any navigation or sensing application that must adapt to operational conditions not fully represented in the pre-training data.

The extension to chain-of-thought reasoning is treated in [78], where the author argues that chain-of-thought prompting enables foundation models to perform multi-step spatial reasoning of the kind relevant to navigation tasks—route planning, landmark identification, spatial-relationship inference—that single-step in-context learning does not support directly. The chain-of-thought angle is particularly relevant to Pillar IV open problem O72 (reasoning-capable AI receivers for complex sensing-and-communication decisions).

In-context learning for GNSS and navigation. The specialization of in-context learning to GNSS and navigation applications is developed in [17], which provides an in-context-learning framework for navigation tasks that adapts to changing signal environments (urban-canyon multipath, ionospheric scintillation, constellation reconfiguration) without retraining. The framework is an operational substrate for the Pillar III open problem O66 (AI-native SAGIN receiver sample complexity) and for the Pillar IV open problem O73 (ICL-based neural-receiver sample-complexity characterization).

Foundation models for geospatial and climate applications. The foundation-model line extends in-context learning to pretrained models for geospatial analysis [79] and for climate modeling [80]. These models adapt the transformer-architecture and pretraining-corpus methodology of natural-language foundation models to the specialized data modalities of geospatial and climate science. The Constellation-Aware Transformer architecture [16], cited in Pillar I (§3.5), is the wireless-relevant instantiation of the same architectural philosophy. Multi-modal foundation models that combine geospatial, atmospheric, and communication signals in a single representation [81] provide a research substrate for the Pillar IV / Pillar III interface open problem O74 (multi-modal SAGIN receivers).

Quantum-ICL for geopotential estimation. The extension of in-context learning to quantum computation for geopotential estimation [82] is a forward-looking bridge to the cross-pillar convergence section (§9), where quantum-enhanced estimation is developed as a cross-pillar frontier. The Pillar IV substrate here is relevant because a quantum-enhanced AI-native receiver is a research direction that Pillar IV needs to treat at the level of capacity bounds; a receiver that can exploit quantum-information resources may achieve Fisher-information-matrix entries that classical receivers cannot.

6.5 Trustworthiness: the provable-safety gap

Pillar IV's treatment of AI safety and trustworthiness, beyond the capacity-theoretic and sample-complexity questions of §6.3, reflects a concern that the current industrial trajectory of AI-native PHY is accumulating specification and prototype artifacts whose trustworthiness has not been established in a form that a safety-critical application can rely on. This subsection develops the concern as a Pillar IV research problem.

The classical wireless physical layer has a well-developed trustworthiness apparatus. A 3GPP-specified decoder has specified performance bounds on specified channel models; the bounds are verifiable; the specifications are subject to regulatory review; and the system as deployed meets performance targets that are within known tolerances of the bounds. A neural receiver trained on synthetic data and deployed in a production network does not naturally carry the same trustworthiness apparatus. Its performance bounds are empirical, established by testing rather than by theorem; its robustness to adversarial inputs is typically uncharacterized; its degradation modes under conditions outside the training distribution are discoverable only by deployment; and its internal operation is in substantial respects opaque to the verification procedures that the classical physical layer admits.

The response of the industrial and research communities to this trustworthiness gap has consisted in substantial part of adversarial-testing and red-teaming protocols: exercising a neural receiver against adversarial inputs designed to elicit failure modes, cataloging the failures, and hardening the system against them. This is a useful practice. The author's critical analysis of this practice [77] argues, however, that red-teaming as the primary trustworthiness methodology is insufficient for applications in which correctness is required to be provable, not merely empirically checked: the set of possible adversarial inputs is in general larger than any red-teaming exercise can cover, and the safety claim “the system passes these red-team tests” is weaker than the safety claim “the system cannot fail in this family of ways, provably.” For applications such as aviation-grade navigation, autonomous-vehicle safety, and medical-device signal processing, the provable-safety claim is the one the application requires, and the provable-safety apparatus for AI-native PHY components has not been constructed.

The Pillar IV research program treats the construction of this apparatus as a central problem (O71) and identifies a specific intermediate target: a safe-fallback specification for AI-native receivers under which the receiver's output can be compared, online, against a classical-receiver output with provable error bounds, and the AI-native output is accepted as the system output only when the comparison passes the provable bounds. This is problem O75 in the subcatalog. The safe-fallback architecture converts the AI-native receiver's opacity from a safety liability into a controlled performance enhancement: the classical receiver provides the safety floor, the AI-native receiver provides the performance ceiling, and the arbitration between them is provable. Figure 6.1 illustrates the autoencoder-as-communication framework together with the safe-fallback architecture overlay.

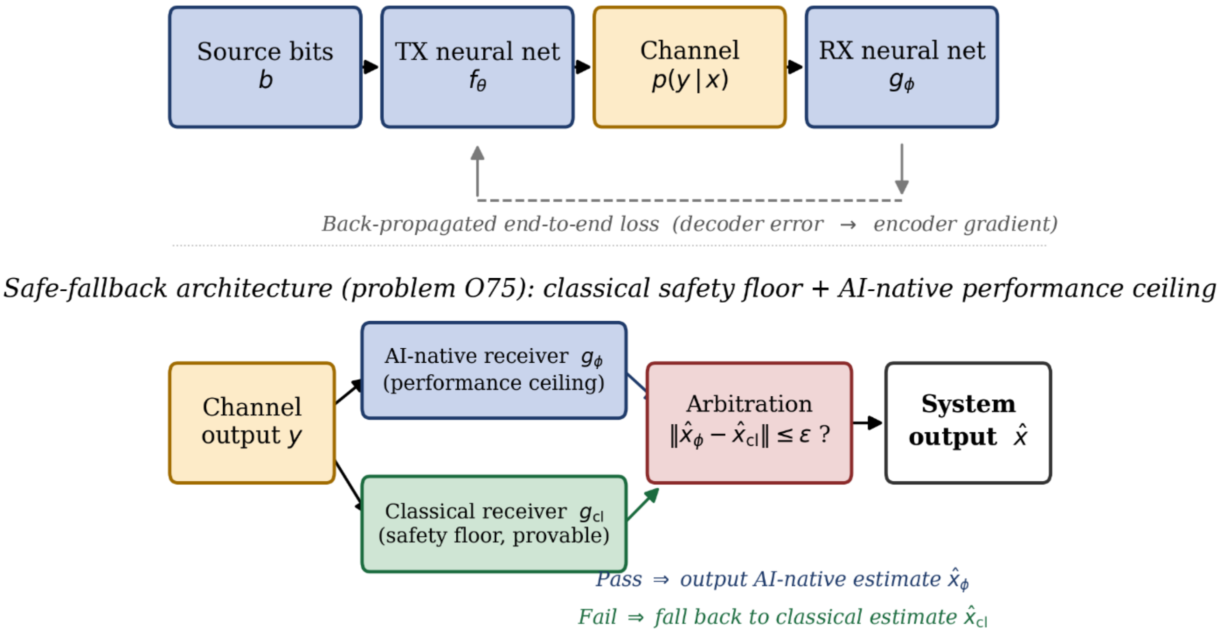

Figure 6.1. Autoencoder-as-communication framework with safe-fallback architecture overlay. (Top) The [69]-[70] end-to-end differentiable chain: transmitter neural network f_θ and receiver neural network g_φ are trained jointly by back-propagating the decoding error through the channel model p(y|x). (Bottom) The safe-fallback architecture of problem O75: the channel output is processed in parallel by the AI-native receiver (performance ceiling) and by a classical receiver (safety floor, provable); an arbitration block compares the two estimates against a provable error bound ε and accepts the AI-native output only when the comparison passes. The architecture converts the AI-native receiver's opacity from a safety liability into a controlled performance enhancement.

6.6 Pillar IV open-problem subcatalog

Sixteen problems are registered in this module, numbered O68 through O83, at the targeted density of 12–15 (the catalog runs slightly over target because the Pillar IV / Pillar III and Pillar IV / Pillar V interfaces deserve explicit entries here).

O68. Neural-codebook capacity gap. Establish the capacity of a communication channel under the restriction that the encoder ranges over codebooks expressible by a neural network of specified architecture and parameter count, and characterize the gap to Shannon capacity as a function of the architecture complexity. Difficulty: Grand Challenge. Prerequisites: Shannon capacity theory, function-approximation theory for neural networks. Effort: 18–30 PhD-months. Spiral potential: defines a Pillar IV sub-literature.

O69. PAC-learning bounds for ISAC neural receivers. Derive sample-complexity bounds for neural receivers in the ISAC setting, specifically: given an ISAC channel at Fisher–Shannon operating point (Rate, FIM^{-1}) and a target gap ε to the CRB, what is the training-sample count N sufficient for a neural receiver of parameter count P to achieve the target with probability $1 - \delta$? This is the PAC-learning specialization of the Pillar I open problem O19. Difficulty: Grand Challenge. Prerequisites: PAC-learning, neural-network approximation, Fisher-information theory. Effort: 18–30 PhD-months. Spiral potential: 5–8 follow-on papers on specific architectures.

O70. End-to-end autoencoder-communication convergence. Characterize the convergence of stochastic gradient descent for end-to-end autoencoder-communication training on canonical channels, producing bounds on the training time and on the probability of converging to a useful local minimum. Difficulty: Hard. Prerequisites: optimization theory, neural-tangent-kernel analysis, autoencoder theory. Effort: 15–20 PhD-months. Spiral potential: 3–4 follow-on papers on specific channel families.

O71. Provable-safety analog for AI-native estimation. Construct a provable-safety specification for AI-native estimators in safety-critical wireless-sensing-and-positioning applications, producing a formal guarantee of the form “the estimator's output lies within ε of the true state with probability at least $1 - \delta$ under operating conditions satisfying [specified assumptions].” Difficulty: Grand Challenge. Prerequisites: neural-network verification theory, Kalman-filter theory, wireless-positioning CRB. Effort: 24–36 PhD-months. Spiral potential: defines a sub-literature at the intersection of AI-safety and wireless-positioning; standards-relevant for aviation, automotive, and medical-device applications.

O72. Chain-of-thought-enabled AI receivers. Characterize the performance gain available to neural receivers that employ chain-of-thought reasoning (multi-step inference) rather than single-step inference, for sensing-and-communication decisions that require multi-step reasoning (e.g., multi-hypothesis channel identification, multi-step routing-and-decoding in SAGIN). Difficulty: Hard. Prerequisites: chain-of-thought theory, sensing-and-communication decision theory. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O73. ICL-based neural-receiver sample complexity. Specialize O69 to in-context-learning-based neural receivers, where the training-sample count is replaced by the context-example count and the performance question is how few context examples are needed to reach within ε of the Bayes-optimal posterior. Difficulty: Hard. Prerequisites: ICL theory, neural-receiver theory. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O74. Multi-modal SAGIN receiver foundation model. Characterize the Fisher–Shannon performance of a multi-modal foundation-model receiver that jointly ingests communication signals, positioning signals, atmospheric-sensing signals, and spatial-context signals in a SAGIN setting. Difficulty: Hard. Prerequisites: multi-modal foundation models, SAGIN capacity-positioning (§5). Effort: 15–20 PhD-months. Spiral potential: 3–4 follow-on papers at the Pillar III / Pillar IV interface.

O75. Safe-fallback architecture for AI-native receivers. Specify a safe-fallback architecture in which an AI-native receiver's output is accepted as the system output only when online comparison against a classical-receiver output satisfies a provable error-bound condition, and characterize the capacity-positioning performance of the combined system as a function of the fallback-threshold setting. Difficulty: Hard. Prerequisites: safe-fallback theory, AI-native and classical receiver theory. Effort: 12–18 PhD-months. Spiral potential: 3–5 follow-on papers on specific safety-critical applications; standards-relevant for safety-critical wireless.

O76. Adversarial robustness of neural ISAC receivers. Characterize the adversarial-robustness properties of neural receivers in ISAC settings, where an adversary may craft perturbations to the sensing target or to the communication signal intended to degrade either the sensing CRB or the communication rate. Difficulty: Hard. Prerequisites: adversarial-robustness theory, ISAC CRB-rate. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers on specific adversarial models.

O77. Out-of-distribution generalization bounds for PHY neural networks. Derive out-of-distribution generalization bounds for neural PHY components, specifically bounds that hold when the deployment channel statistics differ from the training channel statistics in specified ways. Difficulty: Hard. Prerequisites: out-of-distribution generalization theory, channel modeling. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O78. Calibrated uncertainty for neural PHY receivers. Develop calibrated-uncertainty methods for neural PHY receivers whose output uncertainty estimates match the true output error distribution in the sense of calibration metrics (expected calibration error, reliability diagrams), with formal calibration bounds. Difficulty: Hard. Prerequisites: calibration theory, neural-receiver theory. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O79. Learned Zak-domain equalization capacity gap. Continue the Pillar II open problem O41: characterize the Zak-domain capacity achievable by a neural equalizer as a function of its parameter count. Difficulty: Hard. Prerequisites: Zak-OTFS capacity, neural-equalization theory. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O80. Foundation-model compression bounds for CSI feedback. Derive information-theoretic bounds on the CSI-feedback compression achievable by foundation-model-based feedback compressors, in terms of the feedback bit-rate required to reach within ε of the ideal-feedback capacity. Difficulty: Hard. Prerequisites: CSI compression, foundation-model compression theory. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers; 3GPP-relevant.

O81. Energy-efficient neural receivers. Characterize the energy-efficient capacity achievable by neural receivers—that is, the Shannon capacity subject to a constraint on the receiver's computational energy consumption per decoded bit. Difficulty: Hard. Prerequisites: energy-efficient capacity, neural-network computational-complexity analysis. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O82. Transferability of learned PHY components across deployment regimes. Characterize the transferability of learned PHY components (neural receivers, learned codebooks, CSI-feedback autoencoders) across deployment regimes (frequency bands, cell sizes, mobility profiles, hardware families), producing bounds on the transfer-learning sample complexity needed to re-specialize the component to a new regime. Difficulty: Hard. Prerequisites: transfer-learning theory, PHY component theory. Effort: 12–18 PhD-months. Spiral potential: 4 follow-on papers.

O83. Standards-usable performance bounds for AI-native PHY. Develop a general framework under which information-theoretic performance bounds on AI-native PHY components can be expressed in the form standards bodies specify performance requirements (e.g., "block-error rate at least 10^{-3} at SNR X for channel model Y"), with the bounds verifiable by a standards-body-operable procedure. Difficulty: Hard. Prerequisites: 3GPP specification methodology, AI-native PHY theory. Effort: 15–20 PhD-months. Spiral potential: directly relevant to 3GPP Release 20 and 21 normative decisions.

The sixteen problems span the capacity-theoretic foundations (O68), the sample-complexity questions (O69, O73), the training-convergence question (O70), the trustworthiness program (O71, O75, O76, O78), the architectural-reasoning direction (O72), the multi-modal and cross-pillar interfaces (O74, O79, O80), and the realistic-deployment directions (O77, O81, O82, O83). The trustworthiness cluster (O71, O75, O76, O78) is the densest, reflecting the argument in §6.5 that the provable-safety gap is the most consequential foundational problem for safety-critical AI-native-PHY applications.

7. Pillar V — Electromagnetic Information Theory

The fifth pillar treats electromagnetic information theory, the program that aims to characterize the capacity and degrees of freedom of wireless channels under the constraint that the channel is governed by Maxwell's equations rather than by abstract scalar-channel models. The program's distinctive contribution, relative to the classical Shannon-based analysis of wireless channels, is that it treats the electromagnetic field as the underlying information-bearing object and derives capacity bounds that respect the physics of radiation, diffraction, mutual coupling, and near-field structure. The distinction is not ornamental. As antenna apertures become electrically large (holographic MIMO, reconfigurable intelligent surfaces spanning meters or tens of meters) and as operating frequencies extend into regimes where near-field effects persist over commercial communication ranges, the scalar-channel abstraction that served earlier wireless generations loses accuracy, and capacity expressions derived from it become bounds on the wrong object. Electromagnetic information theory replaces the scalar channel with the physical electromagnetic field and produces capacity bounds that hold under Maxwell's equations. The pillar's distinctive opportunity, which this section develops, is that near-field positioning and holographic-aperture design have a substantial overlap with the antenna phase-center-variation calibration theory that the geodesy literature developed for high-accuracy GNSS—a bridge that has not been traversed from either side.

7.1 The foundational EIT line

Electromagnetic information theory in its modern form begins with Migliore's 2008 paper [40], which explicitly posed the question of the relationship between electromagnetics and information theory and argued that the wireless-channel-capacity analysis inherited from Shannon could be productively re-examined under the constraint that the channel's degrees of freedom are bounded by the physics of radiation. Migliore's framework treated the antenna aperture as an electromagnetic object whose capacity-relevant properties are its radiation pattern, its near-field-far-field transition behavior, and its degrees of freedom measured in the Green's-function sense. The framework situated several existing wireless-capacity results as special cases of a more general Maxwell-aware capacity theory.

The degrees-of-freedom angle was developed substantially by Franceschetti, Migliore and Minero in a sequence of papers on the capacity of wireless networks and on the information-theoretic content of radiated electromagnetic fields [83]. Their line established that the degrees of freedom of an electromagnetic field in a specified region, observed from a specified aperture, are bounded above by a geometric quantity computable from the physical dimensions of the aperture and the observation region; the bound generalizes the classical antenna-theory result that a circular aperture of diameter D at wavelength λ supports approximately $(\pi D/\lambda)^2$ degrees of freedom in its radiation pattern, and extends it to arbitrary aperture geometries and observation regions.

The holographic MIMO synthesis of this line is due to Pizzo, Marzetta and Sanguinetti [41], who developed a Fourier plane-wave series expansion for holographic MIMO channels. Their expansion represents the channel as a linear combination of plane waves whose amplitudes are sampled at Nyquist-like intervals in the angular domain, and the expansion provides a compact basis for channel modeling and capacity analysis in the large-aperture regime. The companion paper of Pizzo, Torres, Sanguinetti and Marzetta [84] develops the Nyquist sampling theorem for electromagnetic fields, establishing how densely an aperture must be sampled to avoid aliasing in the angular-domain representation.

The engineering side of the pillar is anchored in the holographic MIMO line developed by Huang, Hu, Alexandropoulos and collaborators [85], and Dardari and Decarli [86] on intelligent surfaces as holographic communication devices. The Gong et al. JSAC treatment of holographic MIMO with arbitrary surface placements and near-field line-of-sight capacity [87] extends the framework to non-planar aperture arrangements and establishes near-field capacity limits for realistic deployment scenarios. The Wang et al. electromagnetic-information-theory synthesis [11] consolidates the research program as of 2024, and the Zhu, Tan, Dai analysis of MIMO capacity and channel estimation for EIT [88] extends the framework to include practical channel-estimation considerations.

Recent work has moved the program toward validated physical prototypes. The Nature Communications paper of Zhang, Cui and collaborators [89] demonstrates a dual-channel near-field holographic MIMO prototype with a programmable digital-coding-metasurface implementation; the paper establishes that the theoretical capacity gains predicted by the near-field HMIMO analysis are experimentally realizable. The near-field communications survey by Liu et al. [90] provides a recent consolidated treatment of the near-field-specific phenomena (spherical-wavefront-induced spatial resolution in beamforming, near-field beam focusing, near-field positioning) that distinguish the holographic MIMO regime from the conventional far-field MIMO regime. The IEEE Information Theory Society's call for papers on electromagnetic information theory [91] confirms the program's legitimacy as an information-theoretic research area.

7.2 HMIMO and near-field theoretical foundations

Holographic MIMO is the engineering instantiation at which EIT's analytical advantages become most visible. An HMIMO aperture is, in its idealized form, a continuous electromagnetic surface capable of emitting and receiving an arbitrary electromagnetic field within the aperture's physical constraints; in practical implementations, it is a dense array of closely-spaced radiating elements (spacing a small fraction of a wavelength) backed by digital control that achieves the continuous-aperture behavior to a desired approximation. The two regimes that HMIMO opens—holographic transmission beyond the conventional far-field and commercial deployments at distances where the near-field extends well beyond the aperture—both require the EIT-based analysis to produce accurate capacity and performance expressions.

The near-field regime is characterized by two distinctive properties. First, the wavefront at the receiver is spherical rather than planar, so that the phase differences across a receive array encode distance information in addition to angle-of-arrival information. This distance information makes positioning-via-received-phase feasible in ways that far-field MIMO does not naturally support, and it introduces a positioning CRB that has a specific structural form relating to the aperture geometry. Second, the field's amplitude across the receive array is range-dependent in a way that differs from the inverse-square-law behavior of far-field propagation; the near-field amplitude profile encodes additional information that the far-field analysis does not capture. Figure 7.1 illustrates the two regimes and the boundary between them.

The capacity analysis of the near-field HMIMO channel was advanced substantially by the Pizzo-Marzetta-Sanguinetti Fourier-plane-wave expansion and by the Gong et al. JSAC treatment of arbitrary surface placements. For a line-of-sight near-field channel between a transmit aperture of specified geometry and a receive aperture of specified geometry, the capacity can be expressed in terms of the eigenvalues of an explicit operator derived from the Green's function of the aperture-to-aperture channel, with the number of significant eigenvalues corresponding to the number of degrees of freedom of the channel. The degrees of freedom in the near-field regime are systematically higher than in the far-field regime for the same aperture sizes and the same operating wavelength, reflecting the fact that the near-field channel carries information in both angular and radial dimensions.

Three foundational questions remain open in the HMIMO and near-field literature. First, the capacity analysis with realistic mutual coupling—the fact that densely-spaced radiating elements exchange electromagnetic energy in ways that modify each element's effective radiation pattern—is incomplete in general form. The idealized uncoupled-aperture assumption used in much of the existing analysis is accurate for element spacings of half a wavelength or larger but becomes increasingly inaccurate in the genuinely holographic regime of sub-quarter-wavelength spacing. Second, the capacity analysis is largely narrowband; wideband HMIMO capacity, in which the aperture's frequency response is also a capacity-relevant object, is substantially open. Third, the Green's-function-based formulation used in the existing analysis assumes idealized boundary conditions (free space, half-space over ground) that do not fully represent realistic deployment environments with multipath reflectors and dielectric obstacles.

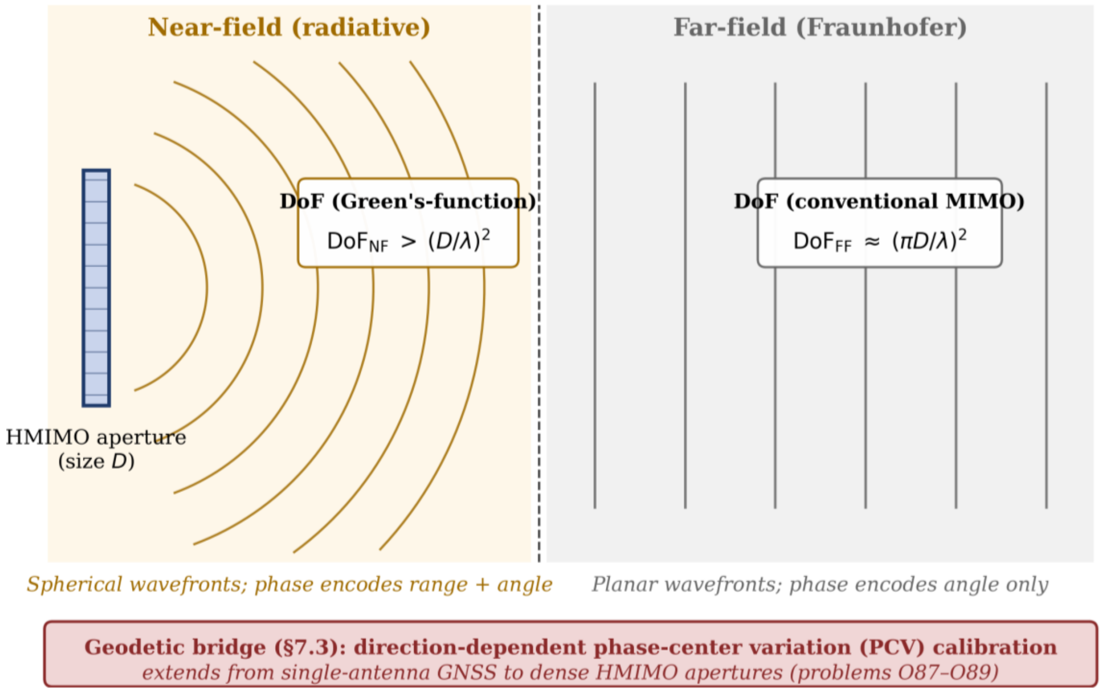

Figure 7.1. Near-field and far-field regimes for HMIMO apertures of size D at wavelength λ. The Fraunhofer distance $d_F = 2D^2/\lambda$ separates the two. In the radiative near-field (left band) the wavefront is spherical and phase across the receive aperture encodes both range and angle; the aperture-to-aperture Green's-function operator supports $\mathrm{DoF\_NF} > (D/\lambda)^2$. In the far-field (right band) the wavefront is planar and phase encodes angle only; $\mathrm{DoF\_FF} \approx (\pi D/\lambda)^2$. The bottom band identifies the geodetic bridge of §7.3: antenna phase-center-variation (PCV) calibration, developed over six decades for single-antenna GNSS, translates with modification to dense HMIMO apertures and anchors problems O87–O89.

7.3 The geodetic bridge: antenna phase-center-variation calibration

The distinctive bridge that Pillar V offers to the geodetic-estimation-theory toolkit is the antenna phase-center-variation (PCV) calibration theory developed for high-accuracy GNSS. This subsection explains the bridge and identifies the specific open problems at which it supplies structure.

The PCV problem in geodesy. An antenna's phase center—the point from which the received signal appears to originate, in the Fresnel sense—is, for any real antenna, not a single geometric point but a direction-dependent and frequency-dependent quantity. For a GNSS antenna required to support millimeter-level positioning in geodetic applications, the direction-dependence and frequency-dependence of the phase center must be characterized to sub-millimeter accuracy and incorporated as a measurement-model term in the receiver's position estimation. The geodesy literature has developed a substantial PCV-calibration apparatus: chamber-based calibration with robot-arm motion to map the phase center over the full sky hemisphere, in-situ calibration via baseline measurements with reference antennas, and field-calibration procedures for operational receivers [92]. The calibration tables produced are incorporated into geodetic receivers' correction models and permit sub-millimeter positioning across multi-frequency GNSS.

The EIT correspondence. The mathematical object that the PCV-calibration apparatus manipulates is, in EIT language, the near-field-to-far-field transformation of the antenna's radiation pattern—equivalently, the aperture-field-distribution-to-radiation-pattern map that the EIT literature analyzes. What geodesy calls "phase-center-variation over the sky" the EIT literature would call "direction-dependent effective phase of the aperture's far-field radiation pattern." The two descriptions are of the same electromagnetic object, expressed in different notational traditions. The geodesy literature has developed the object to the level of operational calibration for six-decade-deployed GNSS receivers; the EIT literature has developed the object to the level of capacity analysis for idealized apertures. The correspondence has not been exploited in either direction.

The near-field positioning bridge. The near-field positioning CRB has the same structural form as the GNSS-positioning CRB extended to include the direction-dependent phase-center contribution: both are inverse-Fisher-information-matrix expressions in which the position-to-measurement Jacobian depends on the antenna's direction-dependent radiation pattern, and both admit DOP-like geometric factorizations. The near-field positioning CRB literature has not, to the author's knowledge, brought in the PCV-calibration machinery from geodesy; the geodetic PCV literature has not, to the author's knowledge, been translated into near-field-positioning terminology. The bridge is exactly the geodetic-wireless isomorphism of §3.3 specialized to Pillar V.

The HMIMO calibration question. The calibration question that HMIMO deployments will face—how to characterize and correct the direction-dependent and range-dependent behavior of a dense sub-wavelength aperture so that positioning and communication performance approach the EIT-theoretic bounds—is, at its mathematical core, a generalized PCV-calibration problem. The geodesy literature's PCV-calibration apparatus was developed for conventional single-antenna GNSS receivers and cannot be applied directly to HMIMO apertures; but its methodology (chamber-based far-field pattern mapping, in-situ calibration via reference sources, field-calibration procedures) translates with modification to the HMIMO setting. This translation is a concrete research program (problems O88, O89 in the subcatalog) that combines the EIT literature's analytical framework with the geodesy literature's operational methodology.

7.4 EIT for ISAC and the atmospheric-interface

The final element of Pillar V's treatment is the interface with Pillars I and III. Two explicit research directions are identified.

EIT-compliant ISAC capacity. The ISAC capacity-distortion theory developed in Pillar I is built on scalar-channel abstractions of the physical wireless channel: the channel is an additive-white-Gaussian-noise link plus specified distortion, and the capacity analysis is performed in the scalar-channel representation. For HMIMO-scale apertures and for commercial-distance deployments where near-field effects persist, the scalar-channel abstraction loses accuracy. An EIT-compliant ISAC capacity analysis—one that characterizes the joint sensing-and-communication capacity of a genuinely electromagnetic channel, with the physics of the channel entering the capacity expression directly—would produce capacity bounds whose accuracy exceeds that of the scalar-channel analyses in the regimes where HMIMO deployments operate. This is problem O90 in the subcatalog and is the Pillar V / Pillar I central interface problem.

Atmospheric structure as EIT degrees of freedom. The ionospheric-scintillation and tropospheric-turbulence effects analyzed in Pillar III (§5.4) as contributions to the Fisher information matrix for positioning have an EIT-side analog: the atmospheric medium through which the electromagnetic field propagates is an additional source of degrees-of-freedom structure, not all of which is captured in the scalar-propagation model. For HMIMO deployments operating over propagation paths long enough that atmospheric effects become non-trivial (aerial, high-altitude, satellite-to-ground), the EIT degrees-of-freedom analysis should include the atmospheric-refractive-index-variation contribution to the channel's capacity structure. This is problem O91 in the subcatalog and is the Pillar V / Pillar III interface problem.

7.5 Pillar V open-problem subcatalog

Twelve problems are registered in this module, numbered O84 through O95, at the targeted density of 10–12.

O84. EIT capacity with realistic mutual coupling. Derive the capacity of an HMIMO aperture in which the radiating elements exchange electromagnetic energy through realistic mutual-coupling interactions, extending the existing uncoupled-aperture analyses to the sub-half-wavelength-spacing regime characteristic of truly holographic apertures. Difficulty: Grand Challenge. Prerequisites: electromagnetic coupling theory, MIMO capacity, Green's-function analysis. Effort: 18–30 PhD-months. Spiral potential: defines a sub-literature on practical HMIMO capacity.

O85. Wideband EIT capacity. Extend the existing narrowband EIT capacity analyses to wideband settings, producing capacity expressions that characterize the aperture's frequency-response structure as an explicit capacity-relevant object. Difficulty: Hard. Prerequisites: wideband capacity theory, EIT Green's-function formalism. Effort: 15–20 PhD-months. Spiral potential: 3–4 follow-on papers.

O86. Near-field capacity in realistic multipath environments. Extend the existing line-of-sight near-field HMIMO capacity to multipath environments, producing capacity expressions that characterize the near-field-plus-multipath regime that will be typical of commercial HMIMO deployments. Difficulty: Hard. Prerequisites: near-field HMIMO, multipath capacity analysis. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O87. Near-field positioning CRB with phase-center-variation calibration. Derive the near-field positioning CRB that explicitly includes the direction-dependent phase-center contribution of the receive aperture, producing bounds that can be compared against the uncalibrated baseline. Difficulty: Hard. Prerequisites: near-field positioning theory, PCV-calibration theory from geodesy. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers on specific aperture geometries.

O88. HMIMO calibration theory. Develop a calibration theory for dense-aperture HMIMO systems analogous to the geodetic PCV-calibration apparatus, with chamber-based, in-situ, and field-calibration procedures specified formally and with calibration-residual-error contributions to the EIT capacity bound characterized. Difficulty: Hard. Prerequisites: geodetic PCV calibration, HMIMO, EIT. Effort: 15–20 PhD-months. Spiral potential: 3–5 follow-on papers on specific calibration procedures; practically impactful for HMIMO product realizations.

O89. PCV-inherited capacity-loss bounds for HMIMO. Quantify the capacity loss from residual PCV calibration error in an HMIMO deployment, as a function of calibration methodology and calibration-residual-error statistics. Difficulty: Hard. Prerequisites: HMIMO capacity, PCV-calibration-residual statistics. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O90. EIT-compliant ISAC capacity. Characterize the joint sensing-and-communication capacity of an HMIMO-scale aperture under an EIT-compliant channel model, producing CRB-rate region expressions that are valid in the regimes where the scalar-channel abstraction used in Pillar I loses accuracy. Difficulty: Grand Challenge. Prerequisites: EIT capacity, ISAC CRB-rate region (§3). Effort: 18–30 PhD-months. Spiral potential: defines the Pillar V / Pillar I interface sub-literature.

O91. Atmospheric contributions to EIT degrees of freedom. Characterize the atmospheric-medium contribution to the degrees-of-freedom structure of an HMIMO channel over aerial, high-altitude, and satellite-to-ground propagation paths, extending the free-space-Green's-function analyses to include refractive-index-variation effects. Difficulty: Hard. Prerequisites: atmospheric propagation, EIT degrees-of-freedom theory. Effort: 12–18 PhD-months. Spiral potential: 2–3 follow-on papers; Pillar V / Pillar III interface problem.

O92. Near-field beam focusing CRB. Derive the CRB on three-dimensional position estimation from near-field beam focusing, producing bounds that relate aperture size, operating wavelength, target range, and estimation accuracy. Difficulty: Accessible–Hard. Prerequisites: near-field beam focusing, estimation-theoretic CRB. Effort: 8–15 PhD-months. Spiral potential: 2–3 follow-on papers on specific aperture geometries.

O93. Capacity of reconfigurable-intelligent-surface systems under EIT constraints. Extend the existing RIS capacity analyses to EIT-compliant formulations where the RIS's reflection pattern is constrained by the physics of a passive electromagnetic surface rather than by an arbitrary reflection-coefficient abstraction. Difficulty: Hard. Prerequisites: RIS capacity theory, EIT. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers.

O94. Metamaterial-based HMIMO capacity. Characterize the capacity of HMIMO implementations based on programmable metamaterials (digital coding metasurfaces), extending the analytical capacity expressions to account for the specific physical constraints of metamaterial implementations. Difficulty: Hard. Prerequisites: metamaterial physics, HMIMO capacity. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O95. Validated prototypes to theory bridges. Systematize the reconciliation between theoretical HMIMO capacity predictions and experimental measurements from validated prototypes (e.g., the Zhang-Cui 2025 Nature Communications prototype), identifying which theoretical idealizations account for the observed gaps and producing refined theoretical predictions that match the experimental results to within specified tolerances. Difficulty: Hard. Prerequisites: HMIMO capacity, experimental validation methodology. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

The twelve problems span the EIT-capacity foundational questions (O84, O85, O86), the geodetic-bridge problems (O87, O88, O89), the cross-pillar interface problems (O90 for Pillar I, O91 for Pillar III), the specific HMIMO-architecture directions (O92, O93, O94), and the theory-to-experiment reconciliation (O95). The geodetic-bridge cluster (O87–O89) is the pillar's distinctive contribution and is the cluster at which the GNSS–wireless isomorphism of §3.3 pays off for Pillar V.

8. Synthesis — The Fisher–Shannon Unification

Sections 3 through 7 developed the five pillars of beyond-6G foundational theory as separate research programs, each with its own literature, its own open problems, and its own connection to the geodetic-estimation-theory toolkit. This section consolidates the five into a single mathematical framework. The consolidation rests on the Fisher–Shannon metric space $\mathscr{F}$_FS introduced in §3.4, which was defined there for the Pillar I setting and which has appeared in specialized forms across Pillars II, III, IV and V without being formalized at the cross-pillar level. The synthesis below treats $\mathscr{F}$_FS as the unifying object of the survey, states the five pillars' specializations of it compactly, and identifies the concept-pillar matrix against which the remaining open problems of the survey are organized.

8.1 The Fisher–Shannon metric space across pillars

The Fisher–Shannon metric space, as introduced in §3.4, is the set of (Rate, FIM^{-1}) pairs attainable as the input distribution ranges over admissible choices, with the matrix-valued Fisher-information inverse treated directly rather than scalarized. Each pillar engages this object in a specific way.

Pillar I (ISAC Information Theory). The Pillar I specialization is the canonical setting: the sensing task is state estimation for a scalar or low-dimensional parameter vector, the communication task is rate maximization on the same channel, and FIM^{-1} decomposes into a signal-strength factor and a geometry factor in the manner of §3.3. The Xiong-Liu 2023 CRB–Rate region is the scalar-CRB projection of $\mathscr{F}$_FS. The five conjectures FS-1 through FS-5 stated in §3.4 are formulated at the Pillar I level but are intended to admit pillar-specific specializations.

Pillar II (Delay-Doppler Capacity Theory). The Pillar II specialization projects FIM^{-1} along the delay-Doppler axes. For a sparse-scatterer delay-Doppler channel with K scatterers, the Fisher information matrix has a block structure in which each scatterer's delay, Doppler, and complex amplitude contribute block entries. The Zak-domain input distribution that optimizes the rate-versus-sensing-FIM tradeoff produces a specific Pareto surface in $\mathscr{F}$_FS whose characterization is the Pillar II analog of the Pillar I CRB–Rate region. The Lampel equivalence of §4.3 is, in Fisher–Shannon language, the statement that the scalar-rate projection of $\mathscr{F}$_FS is representation-invariant between the delay-Doppler and time-frequency bases under specific channel-model assumptions; the full $\mathscr{F}$_FS projection is not representation-invariant, because different bases produce different FIM^{-1} structures for the sensing task.

Pillar III (SAGIN Positioning-Communication Integration). The Pillar III specialization extends the geometry factor of FIM^{-1} across multiple altitude layers and across multiple constellations. The altitude-weighted GDOP introduced in §5.3 is the specific scalarization of the position-projection of FIM^{-1} for the SAGIN receiver. The atmospheric-delay entries of FIM^{-1}, developed in §5.4, are additional block structure that the Pillar III specialization carries and that the Pillar I base formulation does not. The Conjecture FS-5 of §3.4 (GNSS-DOP recovery) is the central SAGIN-relevant assertion about $\mathscr{F}$_FS, and the Pillar III open problem O50 is the full-generality characterization of $\mathscr{F}$_FS in the SAGIN setting.

Pillar IV (AI-Native Physical Layer Theory). The Pillar IV specialization replaces the classical FIM^{-1} with a neural-posterior-based approximation $\hat{\Sigma}$_θ, where $\hat{\Sigma}$_θ is the empirical error covariance of a learned estimator trained on samples from the channel. The gap between FIM^{-1} and $\hat{\Sigma}$_θ is the central Pillar IV object; the PAC-learning bound of O69 characterizes this gap as a function of neural-network parameter count and training-sample size. The neural-codebook capacity gap of O68 is the Rate-component analog: the communication-rate axis of $\mathscr{F}$_FS is reduced by the restriction of the encoder to a learned codebook family. The provable-safety program of O71 and O75 is the question of under what conditions the learned estimator's approximation $\hat{\Sigma}$_θ can be certified to approximate FIM^{-1} within specified tolerances.

Pillar V (Electromagnetic Information Theory). The Pillar V specialization replaces the scalar-channel abstraction implicit in the Pillar I and II formulations with an electromagnetic-field-theoretic channel description. The Rate axis of $\mathscr{F}$_FS is the EIT-compliant capacity developed by the Pizzo-Marzetta-Sanguinetti-Gong line; the FIM^{-1} axis picks up entries corresponding to the antenna's direction-dependent phase-center structure and to the aperture-to-aperture Green's-function degrees of freedom. The geodetic PCV-calibration bridge of §7.3 enters $\mathscr{F}$_FS as explicit measurement-model terms in the Fisher information matrix for aperture-based positioning.

8.2 Pillar interactions at the framework level

The five specializations are not independent; the pillars interact at the $\mathscr{F}$_FS level in ways that §9 will develop as specific convergence frontiers. This subsection states the interactions at the framework level (Figure 8.1).

Pillar I and Pillar II. An ISAC system whose communication waveform is Zak-OTFS produces a Fisher–Shannon operating point whose Rate axis is governed by the Pillar II delay-Doppler capacity analysis and whose FIM^{-1} axis is governed by the Pillar I CRB–Rate structure. The combined characterization is the CRB–Rate region of a Zak-OTFS ISAC system, which is Pillar II's open problem O38 and the program's P1-6 paper.

Pillar I and Pillar III. An ISAC system operating in a multi-constellation SAGIN setting has its sensing task specialized to multi-constellation positioning, and its Fisher–Shannon operating point inherits both the Pillar I CRB–Rate structure and the Pillar III altitude-weighted GDOP structure. The combined characterization is O50 (SAGIN CRB–Rate region).

Pillar I and Pillar IV. An ISAC system with a neural receiver has its sensing task performed by a learned posterior, and its Fisher–Shannon operating point is shifted from FIM^{-1} to $\hat{\Sigma}\_\theta$ by the neural-estimator's approximation error. The combined characterization is O69 (PAC-learning bounds for ISAC neural receivers).

Pillar I and Pillar V. An ISAC system with an HMIMO-scale aperture has its channel model replaced by an EIT-compliant representation, and its Fisher–Shannon operating point is shifted by the near-field phase-center contributions to FIM^{-1}. The combined characterization is O90 (EIT-compliant ISAC capacity).

Pillar II and Pillar III. A LEO SAGIN channel with large Doppler magnitudes and Doppler rates is a natural Zak-OTFS application domain. The combined characterization is O43 (delay-Doppler capacity under LEO channel models) and O58 (SAGIN capacity under Doppler-rate effects).

Pillar II and Pillar IV. A neural equalizer operating in the Zak domain is a cross-pillar object that combines Pillar II's delay-Doppler capacity framework with Pillar IV's learned-receiver-performance framework. The combined characterization is O41 and O79.

Pillar III and Pillar IV. An AI-native SAGIN receiver—in particular, the in-context-learning-based multi-constellation receiver implied by the author's research line—is a cross-pillar object combining Pillar III's multi-constellation estimation framework with Pillar IV's ICL-based neural-receiver framework. The combined characterization is O66 and O74.

Pillar III and Pillar V. A SAGIN deployment using HMIMO apertures at the ground station combines Pillar III's multi-constellation positioning framework with Pillar V's EIT capacity framework; atmospheric-delay contributions to FIM^{-1} cross-link the two. The combined characterization is O91 (atmospheric contributions to EIT DoF) and its extensions.

Pillar IV and Pillar V. A neural receiver on an HMIMO aperture inherits both the Pillar IV neural-posterior-versus-classical-FIM gap and the Pillar V EIT-compliant channel description. The combined characterization extends O68 and O90 jointly.

These pairwise interactions, enumerated here, are the scaffolding on which §9 will identify the five highest-impact convergence frontiers as objects deserving their own sub-literature treatment.

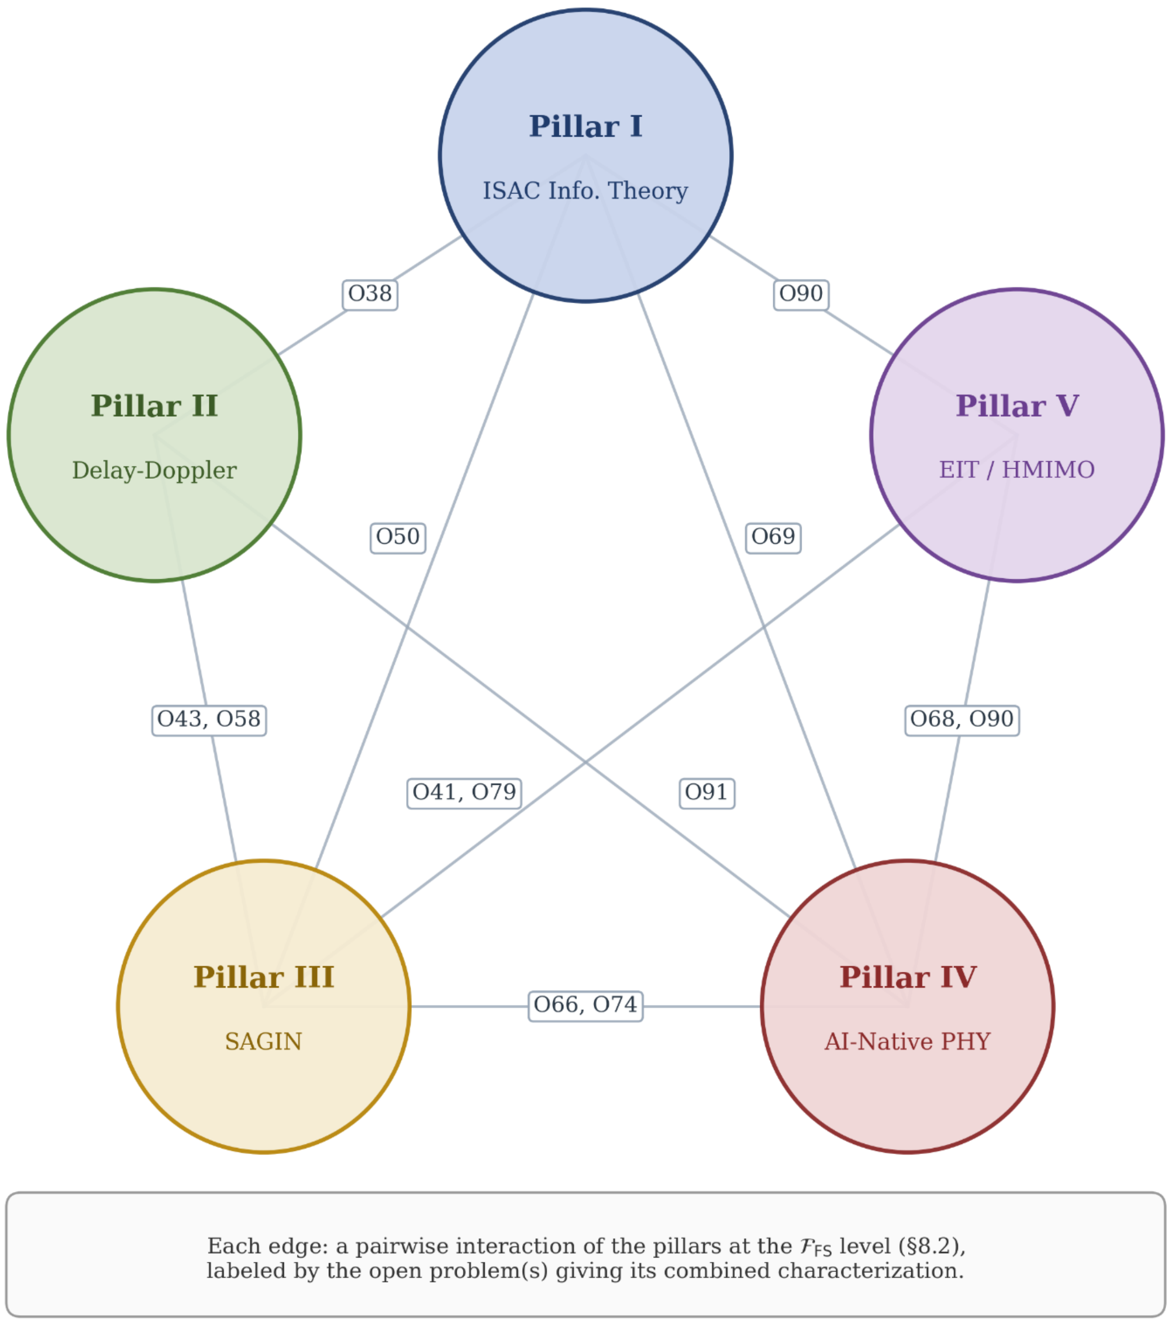

Figure 8.1. Pairwise pillar interactions at the Fisher–Shannon framework level. Each edge links two pillars whose combination is developed in §8.2, labeled by the open problem(s) giving the combined characterization (e.g., O38 for Pillar I × Pillar II). Nine of the ten possible pairs interact substantively at the $\mathscr{F}$_FS level; the five convergence frontiers of Figure 9.1 are the subset of these interactions developed as research frontiers in §9.

8.3 The concept-pillar matrix

The five pillars can be indexed against a set of mathematical concepts drawn from the geodetic-estimation-theory toolkit, producing a concept-pillar matrix that shows, for each concept, which pillars engage it and where in the survey each engagement is developed. Seven concepts are informative at this level of abstraction. Table 8.1 summarizes the matrix; the narrative below develops each concept in turn.

Cramér–Rao bound / Fisher information matrix. Pillar I develops this concept as the sensing-task performance bound; Pillar II specializes it to delay-Doppler scatterer estimation; Pillar III extends it to multi-constellation position-velocity-time estimation; Pillar IV studies its approximation by learned estimators; Pillar V extends it to EIT-compliant aperture-based estimation.

Kalman filter / recursive estimation. Pillar I addresses the temporal completion in Conjecture FS-4 and in open problem O21; Pillar III addresses SAGIN Kalman steady-state in O55; Pillar IV treats the provable-safety aspects in O71 and O75; Pillars II and V touch the concept through their dynamic-channel and dynamic-aperture settings respectively.

Dilution of precision / Fisher-information decomposition. Pillar I introduces this concept as the Pillar I FIM^{-1} decomposition; Pillar III develops the altitude-weighted version; Pillars II, IV and V use it implicitly through the FIM^{-1} projections they analyze.

Ambiguity function / Zak transform. Pillar II is the central pillar for this concept; Pillar I engages it through the ambiguity-function-as-sensing-resolution framing of §3.3; Pillars III and V engage it through the Zak-domain analysis of GNSS-signals and aperture-based waveforms respectively.

Multi-constellation fusion. Pillar III is the central pillar; Pillar I engages it via multi-transmitter ISAC; Pillar IV addresses it through multi-modal foundation-model receivers.

Atmospheric-delay modeling. Pillar III is central; Pillar V engages it through atmospheric contributions to EIT DoF; Pillar I engages it implicitly through multi-path ISAC channel models.

Antenna phase-center-variation calibration. Pillar V is central; Pillars I, II, III and IV engage it through the geodetic-bridge material in §7.3, cross-cited in specific cross-pillar open problems.

Table 8.1. Concept-pillar matrix. For each of seven geodetic-estimation-theory concepts, the table identifies the central pillar (marked **C**) and the pillars that engage the concept substantively (marked ●). Secondary or implicit engagements are marked (○). The matrix is dense: each concept has a central pillar and substantive engagement with at least two other pillars.

Concept	Pillar I	Pillar II	Pillar III	Pillar IV	Pillar V
Cramér–Rao bound / Fisher information matrix	**C**	●	●	●	●
Kalman filter / recursive estimation	**C**	○	●	●	○
Dilution of precision / FIM decomposition	●	○	**C**	○	○
Ambiguity function / Zak transform	●	**C**	○	—	○

Concept	Pillar I	Pillar II	Pillar III	Pillar IV	Pillar V
Multi-constellation fusion	●	—	**C**	●	—
Atmospheric-delay modeling	○	—	**C**	—	●
Antenna phase-center-variation calibration	○	○	●	○	**C**

The concept-pillar matrix is therefore dense: each of the seven concepts has a central pillar and substantive engagement with at least two other pillars. This shows how the toolkit distributes across the partition adopted here, as claimed in §1 that the geodetic-estimation-theory toolkit is the natural mathematical language of beyond-6G foundational theory. The claim is not that a single geodetic concept unlocks all five pillars, but that the toolkit as a whole is distributed across the pillars in a specific, checkable way, with each concept finding its central pillar and its cross-pillar engagements.

8.4 The synthesis open-problem subcatalog

Seven problems are registered in this module, numbered O96 through O102. The problems are of the integrative type: they treat the Fisher–Shannon metric space at the level that cuts across pillar boundaries and produces results relevant to all five.

O96. Full matrix-valued Fisher–Shannon boundary. Prove Conjectures FS-1 and FS-2 (Gaussian-input achievability and convexity) in their general matrix-valued form, not merely in the scalar-CRB projection to which Pillar I's O11 and O12 specialize. Difficulty: Grand Challenge. Prerequisites: all five pillars' pillar-specific $\mathscr{F}$_FS specializations. Effort: 24–36 PhD-months. Spiral potential: defines the unified Fisher–Shannon sub-literature.

O97. Cross-pillar tensorization of $\mathscr{F}$_FS. Characterize how the Fisher–Shannon metric space of a system combining two or more pillars (e.g., a Zak-OTFS ISAC over a SAGIN channel with HMIMO aperture) relates to the single-pillar regions of each component. Prove or disprove a tensorization-type theorem analogous to Conjecture FS-3 for MAC but extended to pillar-combinations. Difficulty: Grand Challenge. Prerequisites: all five pillars. Effort: 24–36 PhD-months. Spiral potential: 5–10 follow-on papers on specific pillar combinations.

O98. Geometry factor universality. Establish that the geometry factor that appears in the GNSS-DOP decomposition, in the Pillar II delay-Doppler-scatterer-geometry decomposition, in the Pillar III SAGIN-altitude decomposition, and in the Pillar V aperture-Green's-function decomposition is, in each case, the same mathematical object up to a pillar-specific specialization—making the "DOP analog" of each pillar a precise theorem rather than a suggestive analogy. Difficulty: Hard. Prerequisites: all five pillars. Effort: 15–20 PhD-months. Spiral potential: 3–4 follow-on papers.

O99. Kalman–Fisher–Shannon trilogy. Derive a unified steady-state-covariance theorem for a Kalman filter operating on measurements drawn from any Fisher–Shannon operating point across the five pillars, generalizing the Pillar I Conjecture FS-4 and the Pillar III open problem

O55 to a cross-pillar result. Difficulty: Hard. Prerequisites: Kalman-filter theory, $\mathscr{F}$_FS across pillars. Effort: 15–20 PhD-months. Spiral potential: foundational for the Kalman-ISAC and Kalman-SAGIN program papers.

O100. Neural approximation of $\mathscr{F}$_FS. Characterize the gap between the classical Fisher–Shannon metric space $\mathscr{F}$_FS and the neural-posterior-based approximation $\mathscr{F}$_FS across all five pillars, producing bounds that apply uniformly. This is the Pillar IV specialization generalized. Difficulty: Grand Challenge. Prerequisites: PAC-learning, all five pillars. Effort: 24–36 PhD-months. Spiral potential: defines a cross-cutting AI-native foundational sub-literature.

O101. Concept-pillar-matrix completeness. Prove or disprove that the seven concepts of §8.3 are a complete basis for the cross-pillar $\mathscr{F}$_FS framework—that is, that any $\mathscr{F}$_FS-related quantity arising in any pillar admits a decomposition whose components are expressible in the seven concepts. Difficulty: Hard. Prerequisites: all pillars, information-geometric decomposition theory. Effort: 18–24 PhD-months. Spiral potential: 3–4 follow-on papers on specific decomposition refinements.

O102. Representation invariance and non-invariance of $\mathscr{F}$_FS. Characterize which projections of $\mathscr{F}$_FS are representation-invariant under basis changes (extending the Lampel equivalence) and which are not. This is the cross-pillar analog of the Pillar II discussion of §4.3. Difficulty: Hard. Prerequisites: all five pillars, representation-theory of channel bases. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

The seven problems span the foundational Fisher–Shannon theorems (O96), the cross-pillar tensorization question (O97, O102), the geometry-factor universality claim (O98), the temporal-completion unification (O99), the neural-approximation unification (O100), and the concept-basis completeness question (O101). They are positioned as the integrative problems that pay off only after all five pillars have been developed individually; they are accordingly placed in this synthesis module rather than in any single pillar.

9. Cross-Pillar Convergence Frontiers

This section identifies the specific regions of the five-pillar taxonomy at which the pillars do not merely touch—as they do through the pairwise interactions enumerated in §8.2—but converge into research frontiers that are distinctive in their own right. Five such frontiers are developed in this module. Each frontier is characterized by three features: it draws on at least two pillars in essential ways (a result restricted to a single pillar is not a cross-pillar frontier); it produces open problems that are not natural within any single pillar and that would not be catalogued in a pillar-specific treatment; and it is expected to generate research output over the 2026–2035 decade at a volume that justifies its identification as a frontier rather than as a single problem. The five frontiers are developed in the subsections below in the order in which they are likely to produce their earliest results, with Frontier 1 (ISAC plus delay-Doppler) the most immediately actionable and Frontier 5 (RIS plus ISAC plus SAGIN) the longest-horizon. Figure 9.1 summarizes the frontiers in a five-pillar layout.

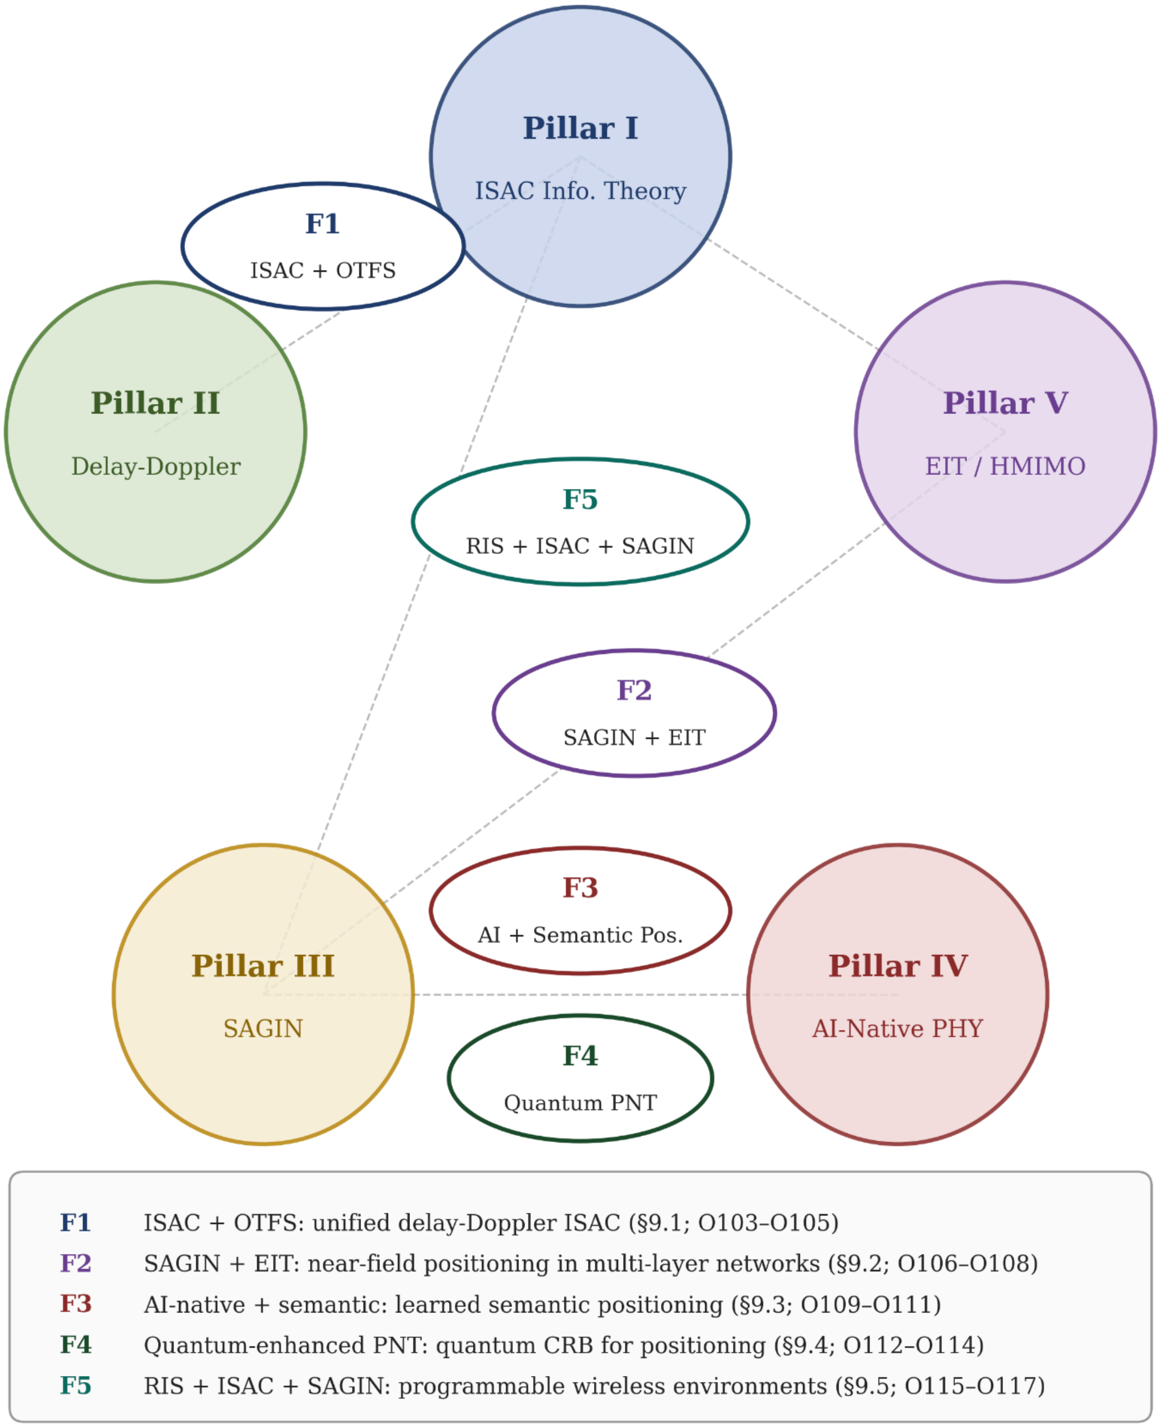

Figure 9.1. Cross-pillar convergence frontiers F1–F5 shown in a five-pillar layout. Each frontier is placed on the edge or at the intersection of the pillars it draws on in essential ways. F1 (ISAC + OTFS) spans Pillars I and II. F2 (SAGIN + EIT) spans Pillars III and V. F3 (AI-native + semantic positioning) spans Pillars IV and III. F4 (Quantum-enhanced PNT) spans Pillars IV and III with a quantum-information extension. F5 (RIS + ISAC + SAGIN) is centrally positioned, drawing on Pillars I, III and V with touch-points in II and IV. The dashed connectors mark which pillars each frontier draws on; the legend at the bottom names the frontier and its anchor open-problem range.

9.1 Frontier 1 — ISAC + OTFS: Unified delay-Doppler ISAC

The first convergence frontier is the unification of Pillar I's ISAC information theory with Pillar II's delay-Doppler capacity theory, producing a systematic treatment of ISAC systems whose waveform is natively delay-Doppler rather than time-frequency. The frontier's central object is the Fisher–Shannon metric space $\mathscr{F}$_FS projected onto delay-Doppler axes and characterized for Zak-OTFS waveforms, with the dual requirement that the communication rate saturate the Pillar II delay-Doppler capacity and that the sensing Fisher information matrix saturate the Pillar I CRB bound for the delay-Doppler scatterer estimation task. The program's paper P1-6 (Zak-OTFS ISAC CRB–Rate region) is the first entry in this frontier; Pillar II's open problem O38 and its extensions are the concrete research agenda.

The frontier's distinctive research opportunity is that Zak-OTFS's embedded-pilot structures, in which a portion of the Zak-domain symbols serves simultaneously as channel-estimation pilot and as sensing waveform, admit Fisher–Shannon analyses that are not available in the time-frequency-pilot equivalents. The Zak-domain pilot structure has specific ambiguity-function properties that determine the sensing CRB directly; the Zak-domain communication symbols have specific capacity-achieving distributions that may or may not coexist with optimal sensing-pilot structure. The tradeoff between these two is not present in the Pillar I or Pillar II treatments in isolation.

Three open problems are characteristic of Frontier 1.

O103. Capacity-achieving Zak-domain pilot density. For a Zak-OTFS ISAC system operating at a target CRB–Rate point, derive the capacity-achieving Zak-domain pilot density as a function of the scatterer-count distribution, the per-scatterer Doppler distribution, and the target sensing-accuracy specification. Difficulty: Hard. Prerequisites: Zak-OTFS, Pillar I and II specializations of $\mathscr{F}$_FS. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers on specific pilot-design optimizations.

O104. Frontier 1 capacity under hardware impairments. Extend the Pillar I + Pillar II CRB–Rate analysis to include the hardware impairments (quantization, phase noise, amplifier nonlinearity) that dominate practical Zak-OTFS-ISAC implementations. Difficulty: Hard. Prerequisites: Frontier 1 capacity, hardware-impaired signal processing. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O105. Multi-user Frontier 1. Extend Frontier 1 to multi-user settings (MAC-ISAC-OTFS, BC-ISAC-OTFS), producing joint capacity-region-and-CRB-region characterizations. Difficulty: Hard. Prerequisites: Frontier 1 single-user, multi-user Pillar I extensions. Effort: 15–20 PhD-months. Spiral potential: 3–4 follow-on papers.

9.2 Frontier 2 — SAGIN + EIT: Near-field positioning in multi-layer networks

The second convergence frontier combines Pillar III's SAGIN multi-constellation positioning framework with Pillar V's electromagnetic-information-theoretic treatment of near-field and holographic apertures. The frontier's central object is the altitude-weighted GDOP computation extended to include the direction-dependent and range-dependent aperture-based contributions to the Fisher information matrix. For SAGIN deployments in which the ground-based receiver uses an HMIMO aperture, the positioning measurements carry the near-field phase-structure information characterized in Pillar V, and the positioning CRB is accordingly reduced relative to the conventional far-field single-antenna receiver assumed in the Pillar III multi-constellation analysis.

The frontier's distinctive research opportunity is that the geodetic phase-center-variation-calibration apparatus, developed in §7.3 as the Pillar V distinctive bridge, is a directly applicable methodology for characterizing the HMIMO aperture's contribution to the SAGIN positioning CRB. The geodetic PCV-calibration procedures were developed for single-antenna receivers, but their methodology (chamber-based pattern mapping, in-situ calibration, field-correction tables) generalizes to HMIMO apertures at the conceptual level; the engineering procedures differ. The information-theoretic outcome is a SAGIN positioning performance characterization that explicitly incorporates the aperture's calibration residual as a term in the combined Fisher information matrix.

Three open problems are characteristic of Frontier 2.

O106. Near-field-HMIMO-augmented altitude-weighted GDOP. Derive the altitude-weighted GDOP formalism (O51) for a SAGIN receiver that uses an HMIMO aperture, explicitly including the near-field phase-center-contribution entries in the Fisher information matrix. Difficulty: Hard. Prerequisites: Pillar III O51, Pillar V §7.3. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers on specific deployment scenarios.

O107. HMIMO SAGIN ground-station calibration. Specify the calibration methodology (chamber-based, in-situ, field-based) for an HMIMO SAGIN ground-station aperture, including the calibration-residual-error contribution to the positioning CRB. Difficulty: Hard. Prerequisites: Pillar V O88, SAGIN receivers. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers; standards-relevant for 3GPP NTN and for LEO-PNT community.

O108. HMIMO-enabled LEO signal-of-opportunity positioning. Characterize the positioning CRB achievable when LEO signals of opportunity are received by an HMIMO ground aperture, exploiting the aperture's angular and near-field-range resolution to extract information that single-antenna receivers cannot. Difficulty: Hard. Prerequisites: Pillar III LEO-PNT, Pillar V near-field. Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers.

9.3 Frontier 3 — AI-native + semantic: Learned semantic positioning

The third convergence frontier combines Pillar IV's AI-native physical-layer theory with the semantic-communication line that Pillar IV's §6.1 touched on through the extension of the autoencoder-as-communication framework to task-defined rather than bit-exact rate metrics. The frontier is specifically concerned with semantic positioning: the application-level task of conveying task-relevant positional information (e.g., “vehicle A is approaching intersection B from the west”) rather than the raw positional estimate itself, with the rate measured in task-relevant units rather than in bits of positional precision.

The frontier's distinctive research opportunity is that semantic positioning admits rate-distortion tradeoffs that exceed the classical Shannon rate for specific task metrics: a system that conveys only task-relevant positional content can do so at lower channel rates than a system that conveys the full positional estimate, with the difference measured by the semantic task's distortion metric. The foundation-model line developed in Pillar IV §6.4 provides concrete substrate: a foundation model that encodes the semantic content of positional data (spatial relationships, intersection topology, navigation-relevant landmarks) can be used as the encoder of a semantic-positioning system whose channel requirements are substantially below those of a full-precision positional-estimate system.

Three open problems are characteristic of Frontier 3.

O109. Semantic-positioning rate-distortion region. Characterize the rate-distortion region for semantic positioning under standard task-metric choices (navigation-relevant, safety-critical, collision-avoidance), producing bounds that show how much rate reduction is available as a function of the task-metric tolerance. Difficulty: Hard. Prerequisites: semantic communication, Pillar III positioning, Pillar IV foundation-model compression (O80). Effort: 12–18 PhD-months. Spiral potential: 3–4 follow-on papers on specific task metrics.

O110. Foundation-model-encoded semantic positioning. Derive the performance of foundation-model encoders for semantic positioning, including the sample-complexity for training the encoder on task-specific data and the generalization bounds for deployment in task conditions outside the training distribution. Difficulty: Hard. Prerequisites: Pillar IV O73 (ICL-based neural receiver) and O80 (foundation-model compression), semantic positioning. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O111. Provable-safety semantic positioning. Adapt the Pillar IV provable-safety program (O71, O75) to semantic-positioning systems, characterizing the conditions under which a semantic-positioning encoder can be certified to preserve safety-critical positional information within specified task-metric tolerances. Difficulty: Grand Challenge. Prerequisites: Pillar IV O71 and O75, semantic positioning. Effort: 18–24 PhD-months. Spiral potential: 3–5 follow-on papers.

9.4 Frontier 4 — Quantum-enhanced PNT: Quantum CRB for positioning

The fourth convergence frontier extends the Fisher–Shannon metric space to quantum-enhanced measurements, producing a quantum Cramér–Rao bound on positioning-and-communication performance that exploits quantum-information resources unavailable to classical measurements. The frontier's central object is the quantum Fisher information matrix $\mathcal{J}$_Q, which generalizes the classical Fisher information matrix to quantum-state-dependent measurements, and its inverse provides the quantum CRB on parameter estimation. For positioning systems that can generate and exploit quantum-correlated probe states, the quantum CRB is, in general, lower than the classical CRB—sometimes substantially so—and the Fisher–Shannon metric space expands correspondingly.

The frontier's distinctive research opportunity has specific author substrate. The author's work on quantum in-context learning for geopotential estimation [82] develops quantum-enhanced estimation in the context of gravity-gradient sensing for geopotential field characterization, producing specific quantum-ICL architectures whose estimation performance exceeds the classical limit. The information-theoretic question that Pillar-level analysis raises is the extension from geopotential estimation to the full PNT (positioning-navigation-timing) task: under what conditions does a quantum-enhanced receiver architecture produce Fisher-information-matrix entries for the PNT state vector that exceed the classical bounds, and by how much?

Three open problems are characteristic of Frontier 4.

O112. Quantum PNT CRB. Derive the quantum CRB on the full PNT state vector (receiver position, velocity, clock offset and frequency) achievable by a quantum-enhanced receiver exploiting quantum-correlated probe states, extending the geopotential-estimation quantum CRB to the PNT setting. Difficulty: Grand Challenge. Prerequisites: quantum Fisher information, GNSS PNT theory. Effort: 18–30 PhD-months. Spiral potential: defines a quantum-PNT sub-literature.

O113. Quantum-ICL for PNT. Extend the quantum-in-context-learning framework of geopotential estimation to the PNT task, producing specific quantum-ICL architectures for positioning receivers and characterizing their sample complexity. Difficulty: Grand Challenge. Prerequisites: Pillar IV ICL theory, quantum computation, PNT theory. Effort: 18–24 PhD-months. Spiral potential: 3–5 follow-on papers.

O114. Quantum-classical hybrid SAGIN. Characterize the achievable capacity-positioning region for a SAGIN in which specific links are quantum-enhanced and others are classical, with hybrid quantum-classical estimation combining across links. Difficulty: Grand Challenge. Prerequisites: Frontier 4 quantum PNT, SAGIN theory. Effort: 18–24 PhD-months. Spiral potential: 3–4 follow-on papers.

9.5 Frontier 5 — RIS + ISAC + SAGIN: Programmable wireless environments

The fifth convergence frontier is the programmable-wireless-environment frontier, in which reconfigurable intelligent surfaces (RIS), deployed across the heterogeneous SAGIN infrastructure and configured adaptively to support joint communication-and-sensing, produce a wireless environment whose channel structure is itself a programmable object. The frontier combines Pillar I (the ISAC capacity-distortion analysis that RIS configurations must support), Pillar III (the multi-layer architecture in which RIS elements are deployed), Pillar V (the electromagnetic-information-theoretic framework that RIS physics is most accurately described in), and implicitly Pillars II and IV through the RIS waveform and AI-driven configuration subfrontiers.

The frontier's distinctive research opportunity is that RIS deployment density and RIS configuration programmability produce a channel whose $\mathscr{F}$_FS is not a fixed set but is a controllable object whose operating point can be driven by the RIS control plane. For SAGIN applications in which channel conditions vary rapidly—high-mobility urban environments, satellite-to-ground links with changing geometry, mixed line-of-sight and blocked-propagation scenarios—RIS-programmable $\mathscr{F}$_FS represents a qualitatively new capability relative to fixed-channel analyses. The capacity-positioning region available to a user in such an environment is the envelope over RIS configurations, and characterizing this envelope is the Frontier 5 central analytical object.

Three open problems are characteristic of Frontier 5.

O115. RIS-programmable $\mathscr{F}$_FS envelope. Characterize the envelope of Fisher–Shannon operating points achievable by programming an RIS deployment of specified size, element count, and control-plane update rate, for a user in a specified geometry. Difficulty: Hard. Prerequisites: RIS capacity theory, Pillar I CRB–Rate, $\mathscr{F}$_FS framework. Effort: 15–20 PhD-months. Spiral potential: 4–5 follow-on papers.

O116. RIS configuration optimization under real-time constraints. Characterize the achievable RIS configuration as a function of control-plane bandwidth, user-state update rate, and online-optimization computational budget. Difficulty: Hard. Prerequisites: RIS real-time control, Pillar IV neural-policy learning. Effort: 12–18 PhD-months. Spiral potential: 3 follow-on papers.

O117. SAGIN-integrated RIS deployment optimization. Characterize the joint deployment-configuration optimization for an RIS population distributed across a multi-layer SAGIN (on terrestrial infrastructure, on HAPS platforms, on LEO satellites), with the objective of maximizing a specified Fisher–Shannon operating point across user populations. Difficulty: Grand Challenge. Prerequisites: Pillar III, RIS deployment, $\mathscr{F}$_FS optimization. Effort: 18–24 PhD-months. Spiral potential: 4 follow-on papers.

9.6 Additional cross-pillar open problems

Beyond the fifteen problems catalogued under Frontiers 1–5 (three per frontier), three additional cross-pillar problems are of sufficient generality or consequence to be listed here rather than under a specific frontier.

O118. Cross-pillar open-problem dependency graph. Produce a formal dependency graph among the open problems of the survey (O1–O117), identifying which problems' solutions are prerequisites to which others', and use the graph to produce a dependency-respecting research-program sequencing. Difficulty: Accessible. Prerequisites: all catalogued problems. Effort: 6–10 PhD-months. Spiral potential: 2 follow-on papers on program-management methodology.

O119. Cross-pillar experimental validation framework. Specify an experimental-validation framework under which claims made across multiple pillars can be tested on a common testbed, producing quantitative cross-pillar-validation-data that can then be used to refine theoretical predictions in the manner of Pillar V's O95 applied across pillars. Difficulty: Hard. Prerequisites: all pillars, experimental wireless methodology. Effort: 18–24 PhD-months. Spiral potential: 3–5 follow-on papers on specific testbed campaigns.

O120. Cross-pillar standards-translation framework. Specify a translation framework from the five-pillar foundational theory to standards-usable performance specifications (3GPP-style, ITU-R-style, WRC-decision-style), generalizing Pillar IV's O83 to the full survey. Difficulty: Hard. Prerequisites: all pillars, 3GPP/ITU-R specification methodology. Effort: 15–20 PhD-months. Spiral potential: directly influential on Release 21 and 22 decision-making.

Eighteen problems are registered across this module, at the targeted density of 15–20 cross-pillar entries. The cumulative program open-problem count now stands at 120, and the survey is complete on the open-problem-catalog side. Frontiers 1–5 contain fifteen of the eighteen new problems, with three problems in the cross-pillar general subsection (O118–O120). Frontier 4 (Quantum-PNT) is uniformly Grand Challenge; Frontiers 1, 2, and 3 are predominantly Hard with specific Grand Challenge entries (O111 in Frontier 3); Frontier 5 has a Grand Challenge in O117. The distribution reflects the horizons: Frontier 4 is earliest-research and hardest; Frontier 1 is most immediately actionable.

10. Master Open-Problem Catalog

Modules 1 through 9 registered 120 open problems as they developed the historical-pattern argument (Module 1), the diagnosis of the 6G theoretical landscape (Module 2), the five pillars (Modules 3–7), the Fisher–Shannon synthesis (Module 8), and the cross-pillar convergence frontiers (Module 9). This section consolidates those 120 problems into a single numbered catalog with consistent classification metadata, organized first for navigability and second for research-program sequencing. The section has four subsections: §10.1 states the classification rubric, §10.2 presents the consolidated problem list organized by pillar and by difficulty, §10.3 provides the master classification table, and §10.4 presents the cross-reference table mapping problems to the planned papers of the 7G Foundational Research Program.

10.1 Classification rubric

Every problem in the catalog is classified along four dimensions.

Dimension 1 — Difficulty. Three levels: *Accessible* — a problem that a competent PhD student with the stated prerequisites can complete within a standard doctoral timeline as a first chapter or as a substantial conference paper. Accessible problems typically admit known-method solutions that have not been applied to the specific setting, or they require only moderate extension of existing results. *Hard* — a problem whose solution requires non-trivial new analytical or methodological work. Hard problems typically consume 12–24 PhD-months and produce a journal paper of substantial novelty. Hard is the default difficulty level of the catalog; the majority of problems are Hard. *Grand Challenge* — a problem whose solution is expected to require coordinated effort across multiple research groups, extended research programs, or novel methodological developments. Grand Challenge problems typically consume multiple PhD-student-years and produce seminal results. Grand Challenge problems are intended as research-program organizing objects, not as single-paper targets.

Dimension 2 — Prerequisite knowledge. Each problem carries a short list of the research areas in which the investigator needs working competence. Prerequisites are stated at the level of textbook-chapter granularity rather than at the level of individual papers.

Dimension 3 — Estimated effort in PhD-months. The estimate is the number of doctoral-student months of full-time research effort plausibly required to produce a publication-quality solution, including literature work, analytical development, possible simulation or experiment, and paper-writing. The provenance of these estimates should be stated plainly, because they are presented in tabular form and may otherwise be read as measurements. They are the author's judgments, anchored on three reference points: the effort required by published results of comparable scope in the same subfields, the standard doctoral timeline in wireless communications and geodesy, and the depth of the prerequisite list each problem carries. They have not been validated against completed work on these specific problems, for the sufficient reason that the problems are open; no second assessor has scored them; and no inter-rater calibration has been performed. They should be read as order-of-magnitude guidance for portfolio selection rather than as predictions, and the same caveat applies with greater force to the spiral-potential figures of Dimension 4, which forecast a publication response that depends on community uptake as much as on the problem itself.

Dimension 4 — Spiral potential. The spiral-potential metric estimates the number of follow-on research papers the problem's solution is expected to enable within a five-year horizon. High-spiral-potential problems are those whose solution opens up specific subquestions whose pursuit is publishable; low-spiral-potential problems are those whose solution is largely self-contained. The metric is intended for researchers and research-program architects making portfolio decisions.

The four dimensions are not independent. Accessible problems typically have low-to-moderate spiral potential; Grand Challenge problems typically have high spiral potential and broad prerequisite lists. The classification is useful principally for the following two planning questions: (i) given a specific researcher or group with specified prerequisites, which catalog problems are tractable? and (ii) given a portfolio choice among problems, which combination produces the largest expected research-output stream?

10.2 Consolidated problem list

The problems are grouped by pillar and sub-grouped by difficulty within each pillar. Where a problem spans pillars (cross-pillar open problems registered in Module 9 or Module 8), it is listed under the primary pillar identified at its introduction, with a cross-reference note.

Module 1 — Introduction & framing (4 problems, O1–O4). Grand Challenge: O1 (Successor theorem to Shannon/Telatar/Arıkan). Hard: O2 (2026–2030 standards-window responsiveness), O4 (Foundational vs architectural taxonomy). Accessible: O3 (Geodesy-wireless citation-graph bibliometrics).

Module 2 — State of 6G (5 problems, O5–O9). Hard: O5 (IMT-2030 targets to IT-ET bounds), O6 (3GPP AI/ML-PHY provable-bound framework), O9 (Historical-gap framework).

Accessible: O7 (Usage-scenario-to-pillar taxonomy), O8 (Foundational-theorem lead-time empirical).

Pillar I (Module 3) — ISAC Information Theory (22 problems, O10–O31). Grand Challenge: O14 (Broadcast-channel ISAC boundary), O15 (Interference-channel ISAC), O19 (Neural-receiver ISAC sample complexity). Hard: O10, O11, O12, O13, O16, O17, O18, O20, O21, O22, O23, O24, O25, O27, O29, O30, O31 (17 problems). Accessible or Accessible–Hard: O26 (Fading-channel ISAC capacity — varies by model), O28 (ISAC under hardware impairments).

Pillar II (Module 4) — Delay-Doppler Capacity Theory (18 problems, O32–O49). Grand Challenge: O32 (Native sparse delay-Doppler capacity theorem). Hard: O33, O34, O35, O36, O37, O38, O40, O41, O43, O44, O45, O46, O48, O49 (14 problems). Accessible or Accessible–Hard: O39 (Zak-domain ambiguity-function synthesis), O42 (Crystallization under hardware constraints), O47 (OTFS-vs-Zak-OTFS capacity gap).

Pillar III (Module 5) — SAGIN Positioning-Communication (18 problems, O50–O67). Grand Challenge: O50 (Multi-constellation SAGIN CRB–Rate region). Hard: O51, O52, O53, O55, O56, O57, O58, O59, O60, O61, O62, O63, O64, O65, O66, O67 (16 problems). Accessible or Accessible–Hard: O54 (Inter-constellation timing-offset implications).

Pillar IV (Module 6) — AI-Native Physical Layer (16 problems, O68–O83). Grand Challenge: O68 (Neural-codebook capacity gap), O69 (PAC-learning bounds for ISAC neural receivers), O71 (Provable-safety analog for AI-native estimation). Hard: O70, O72, O73, O74, O75, O76, O77, O78, O79, O80, O81, O82, O83 (13 problems). Accessible: none. Pillar IV's problems are uniformly Hard or Grand Challenge, reflecting the pillar's foundational immaturity.

Pillar V (Module 7) — Electromagnetic Information Theory (12 problems, O84–O95). Grand Challenge: O84 (EIT capacity with realistic mutual coupling), O90 (EIT-compliant ISAC capacity). Hard: O85, O86, O87, O88, O89, O91, O93, O94, O95 (9 problems). Accessible or Accessible–Hard: O92 (Near-field beam-focusing CRB).

Synthesis (Module 8) — Fisher–Shannon Unification (7 problems, O96–O102). Grand Challenge: O96 (Full matrix-valued FS boundary), O97 (Cross-pillar tensorization), O100 (Neural approximation of $\mathscr{F}$_FS). Hard: O98, O99, O101, O102 (4 problems).

Cross-pillar convergence (Module 9) — Five Frontiers plus general (18 problems, O103–O120). Grand Challenge: O111 (Provable-safety semantic positioning), O112 (Quantum PNT CRB), O113 (Quantum-ICL for PNT), O114 (Quantum-classical hybrid SAGIN), O117 (SAGIN-integrated RIS deployment optimization). Hard: O103, O104, O105, O106, O107, O108, O109, O110, O115, O116, O119, O120 (12 problems). Accessible: O118 (Cross-pillar open-problem dependency graph).

Aggregate statistics.

Total problems: 120. Grand Challenge: 19 (16%). Hard: 90 (75%). Accessible or Accessible–Hard: 11 (9%). Average estimated effort: approximately 17 PhD-months per problem, computed as the mean of the range midpoints recorded in Table 10.1 across the 118 problems carrying numeric estimates (O1 is open-ended and O26 is stated per model); if all 120 problems were pursued sequentially by a single researcher, the total effort would approach 2,000 PhD-months, or roughly 165 PhD-years. The catalog is explicitly designed to be attacked in parallel by a research community rather than sequentially by any single group.

Coverage of the geodetic-toolkit concepts. Measured against the seven-concept basis of §8.3: Cramér–Rao bound / Fisher information is engaged by problems across all five pillars (O10, O21, O38, O50, O68, O69, O87, O90, and extensions—total of ≈35 problems directly engage this concept). Kalman filtering is engaged by O21, O22, O55, O99 and extensions (≈15 problems). Dilution of precision and Fisher-information decomposition is engaged by the DOP-like problems across pillars (O51, O98, and extensions, ≈10 problems). Ambiguity function / Zak transform is primarily Pillar II (O33, O34, O39, O46) with cross-pillar engagements. Multi-constellation fusion is Pillar III central (O50, O52, O54). Atmospheric delay is Pillar III central with Pillar V extension (O59, O60, O61, O91). Antenna phase-center calibration is Pillar V central with cross-pillar extensions (O87, O88, O89, O106, O107). The distribution of concepts across problems is broad, illustrating the §8.3 observation that the geodetic toolkit distributes across pillars rather than concentrating in one.

10.3 Master classification table

The following table indexes all 120 problems by (ID, title, primary pillar, difficulty, effort estimate in PhD-months, spiral-potential estimate). The table is intended as a navigable reference object that readers of the survey can consult to identify problems matching their interests, prerequisites, and research-program horizons. Problem titles are abbreviated to single-line headers; the full problem statements appear in the cited modules. Effort estimates are expressed as ranges where the module-level catalog gave a range.

Table 10.1. Master open-problem classification: 120 problems (O1–O120) indexed by primary pillar, difficulty, estimated PhD-months of effort, and spiral potential (expected number of follow-on papers). The table is designed to flow across pages; header row shades repeat on continuation pages under the Word default.

Problem ID	Title	Primary pillar	Difficulty	Effort (PhD-months)	Spiral potential
O1	Successor theorem (Shannon/Telatar/Arıkan)	Intro	Grand Challenge	multi-year	defines subfield
O2	2026–2030 standards window responsiveness	Intro	Hard	12–18	informs programs
O3	Geodesy-wireless citation-graph	Intro	Accessible	4–8	2–3
O4	Foundational vs architectural taxonomy	Intro	Hard	8–12	3–5
O5	IMT-2030 targets to IT-ET bounds	State-of-6G	Hard	18–24	6–10
O6	3GPP AI/ML-PHY provable-bound framework	State-of-6G	Hard	18–30	5–8

Problem ID	Title	Primary pillar	Difficulty	Effort (PhD-months)	Spiral potential
O7	Usage-scenario-to-pillar taxonomy	State-of-6G	Accessible	6–12	3–5
O8	Foundational-theorem lead-time empirical	State-of-6G	Accessible	4–8	2–3
O9	Historical-gap framework	State-of-6G	Hard	12–18	3–5
O10	Full FS boundary MIMO Gaussian ISAC	I	Hard	12–18	3–4
O11	FS-1 proof (Gaussian input achievability)	I	Hard	12–18	2–3
O12	FS-2 proof (convexity)	I	Hard	8–12	2
O13	FS-3 proof (MAC tensorization)	I	Hard	12–18	3–5
O14	Broadcast-channel ISAC boundary	I	Grand Challenge	24–36	sub-literature
O15	Interference-channel ISAC	I	Grand Challenge	24–36	sub-literature
O16	Relay-channel ISAC	I	Hard	18–24	3–5
O17	Extended-target ISAC CRB	I	Hard	12–18	3–4
O18	Wideband ISAC CRB–Rate	I	Hard	12–18	3
O19	Neural-receiver ISAC sample complexity	I	Grand Challenge	18–30	sub-literature
O20	Continuous-waveform ISAC CRB–Rate	I	Hard	12–18	2–3
O21	Kalman-ISAC steady-state (FS-4 proof)	I	Hard	12–18	defines Kalman-ISAC
O22	Kalman-ISAC transient response	I	Hard	12–18	3
O23	JSCC ISAC separation	I	Hard	12–18	2–3
O24	Multi-target ISAC CRB–Rate	I	Hard	18–24	4–5
O25	Privacy-constrained ISAC capacity	I	Hard	12–18	2–3
O26	Fading-channel ISAC capacity	I	Accessible–Hard	8–15 per model	5–7
O27	Finite-blocklength ISAC	I	Hard	15–20	3
O28	ISAC under hardware impairments	I	Accessible–Hard	8–15	4–6
O29	ISAC with imperfect CSI	I	Hard	12–18	3
O30	ISAC energy efficiency	I	Hard	12–18	3
O31	ISAC under secrecy constraints	I	Hard	12–18	3

Problem ID	Title	Primary pillar	Difficulty	Effort (PhD-months)	Spiral potential
O32	Native sparse DD capacity theorem	II	Grand Challenge	24–36	sub-literature
O33	Zak-domain Woodward principle	II	Hard	8–12	2–3
O34	Zak-domain thumbtack pulse classification	II	Hard	8–15	3–5
O35	Water-filling in DD domain	II	Hard	12–18	3
O36	DD finite-blocklength theory	II	Hard	15–20	3
O37	DD MAC capacity region	II	Hard	15–20	3
O38	Zak-OTFS CRB–Rate region	II	Hard	12–18	3–4
O39	Zak-domain ambiguity-function synthesis	II	Accessible	6–12	3–5
O40	DD capacity fractional Doppler	II	Hard	12–18	2–3
O41	Learned Zak-domain equalization	II	Hard	12–18	3
O42	Crystallization under hardware constraints	II	Accessible–Hard	8–12	3–4
O43	DD capacity LEO channel models	II	Hard	12–18	3–4
O44	DD secrecy capacity	II	Hard	12–18	2–3
O45	DD capacity mmWave beamforming	II	Hard	12–18	3–4
O46	Pulse-shape-Zak joint design	II	Hard	12–18	3
O47	OTFS-vs-Zak-OTFS capacity gap	II	Accessible	6–10	2
O48	DD capacity imperfect CSI	II	Hard	12–18	3
O49	DD capacity non-cooperative sensor	II	Hard	12–18	2–3
O50	Multi-constellation SAGIN CRB–Rate	III	Grand Challenge	18–30	sub-literature
O51	Altitude-weighted GDOP formalism	III	Hard	12–18	3–4
O52	LEO-MEO capacity-positioning integration	III	Hard	12–18	3–4
O53	HAPS-terrestrial handover continuity	III	Hard	12–18	3
O54	Inter-constellation timing-offset implications	III	Accessible–Hard	8–12	2–3
O55	SAGIN Kalman steady-state	III	Hard	12–18	3
O56	SAGIN robustness under degradation	III	Hard	12–18	3–4
O57	SoOp information-theoretic characterization	III	Hard	12–18	3
O58	SAGIN capacity Doppler-rate effects	III	Hard	12–18	3
O59	Atmospheric delay as FIM entry	III	Hard	12–18	3–4

Problem ID	Title	Primary pillar	Difficulty	Effort (PhD-months)	Spiral potential
O60	Ionospheric scintillation SAGIN decoding	III	Hard	12–18	3
O61	Multi-frequency SAGIN capacity-positioning	III	Hard	12–18	4–6
O62	Beyond-Earth SAGIN (lunar/cislunar)	III	Hard	18–24	3–5
O63	SAGIN privacy constraints	III	Hard	12–18	2–3
O64	SAGIN spectrum-sharing constraints	III	Hard	12–18	3
O65	Energy-efficient SAGIN capacity-positioning	III	Hard	12–18	3
O66	AI-native SAGIN receiver sample complexity	III	Hard	15–20	3–4
O67	SAGIN capacity under GNSS denial	III	Hard	12–18	3–4
O68	Neural-codebook capacity gap	IV	Grand Challenge	18–30	sub-literature
O69	PAC-learning bounds ISAC neural receivers	IV	Grand Challenge	18–30	5–8
O70	End-to-end autoencoder convergence	IV	Hard	15–20	3–4
O71	Provable-safety analog AI estimation	IV	Grand Challenge	24–36	sub-literature
O72	Chain-of-thought AI receivers	IV	Hard	12–18	3
O73	ICL-based neural-receiver sample complexity	IV	Hard	12–18	3
O74	Multi-modal SAGIN foundation-model receiver	IV	Hard	15–20	3–4
O75	Safe-fallback architecture AI receivers	IV	Hard	12–18	3–5
O76	Adversarial robustness neural ISAC	IV	Hard	12–18	3
O77	Out-of-distribution generalization PHY	IV	Hard	12–18	3
O78	Calibrated uncertainty neural PHY	IV	Hard	12–18	3
O79	Learned Zak-domain capacity gap	IV	Hard	12–18	3
O80	Foundation-model CSI compression bounds	IV	Hard	12–18	3–4
O81	Energy-efficient neural receivers	IV	Hard	12–18	3
O82	Transferability learned PHY components	IV	Hard	12–18	4
O83	Standards-usable AI-native PHY bounds	IV	Hard	15–20	standards-relevant

Problem ID	Title	Primary pillar	Difficulty	Effort (PhD-months)	Spiral potential
O84	EIT capacity realistic mutual coupling	V	Grand Challenge	18–30	sub-literature
O85	Wideband EIT capacity	V	Hard	15–20	3–4
O86	Near-field capacity multipath environments	V	Hard	12–18	3
O87	Near-field positioning CRB with PCV	V	Hard	12–18	3–4
O88	HMIMO calibration theory	V	Hard	15–20	3–5
O89	PCV-inherited HMIMO capacity-loss bounds	V	Hard	12–18	3
O90	EIT-compliant ISAC capacity	V	Grand Challenge	18–30	sub-literature
O91	Atmospheric contributions to EIT DoF	V	Hard	12–18	2–3
O92	Near-field beam focusing CRB	V	Accessible–Hard	8–15	2–3
O93	RIS capacity under EIT constraints	V	Hard	12–18	3–4
O94	Metamaterial HMIMO capacity	V	Hard	12–18	3
O95	Theory-to-experiment HMIMO reconciliation	V	Hard	12–18	3
O96	Full matrix-valued FS boundary	Synthesis	Grand Challenge	24–36	unified FS
O97	Cross-pillar tensorization of FS	Synthesis	Grand Challenge	24–36	5–10
O98	Geometry factor universality	Synthesis	Hard	15–20	3–4
O99	Kalman-FS-trilogy cross-pillar	Synthesis	Hard	15–20	foundational
O100	Neural approximation of FS	Synthesis	Grand Challenge	24–36	sub-literature
O101	Concept-pillar matrix completeness	Synthesis	Hard	18–24	3–4
O102	FS representation-invariance characterization	Synthesis	Hard	12–18	3
O103	Capacity-achieving Zak-domain pilot density	F1	Hard	12–18	3–4
O104	Frontier 1 hardware impairments	F1	Hard	12–18	3
O105	Multi-user Frontier 1	F1	Hard	15–20	3–4
O106	Near-field-HMIMO-augmented GDOP	F2	Hard	12–18	3–4
O107	HMIMO SAGIN ground-station calibration	F2	Hard	12–18	3
O108	HMIMO-enabled LEO SoOp positioning	F2	Hard	12–18	3–4

Problem ID	Title	Primary pillar	Difficulty	Effort (PhD-months)	Spiral potential
O109	Semantic-positioning rate-distortion region	F3	Hard	12–18	3–4
O110	Foundation-model semantic positioning	F3	Hard	12–18	3
O111	Provable-safety semantic positioning	F3	Grand Challenge	18–24	3–5
O112	Quantum PNT CRB	F4	Grand Challenge	18–30	sub-literature
O113	Quantum-ICL for PNT	F4	Grand Challenge	18–24	3–5
O114	Quantum-classical hybrid SAGIN	F4	Grand Challenge	18–24	3–4
O115	RIS-programmable FS envelope	F5	Hard	15–20	4–5
O116	RIS real-time configuration	F5	Hard	12–18	3
O117	SAGIN-integrated RIS deployment	F5	Grand Challenge	18–24	4
O118	Cross-pillar dependency graph	CP-gen	Accessible	6–10	2
O119	Cross-pillar experimental validation	CP-gen	Hard	18–24	3–5
O120	Cross-pillar standards-translation	CP-gen	Hard	15–20	standards-relevant

10.4 Cross-reference: problems to program papers

The 7G Foundational Research Program has planned a sequence of papers in Phases 0, 1, and beyond. This subsection identifies, for each planned Phase-0 and Phase-1 paper, the subset of catalog problems that the paper addresses directly.

Table 10.2. Program paper to catalog-problem cross-reference. Of the 120 catalog problems, approximately 25 are directly addressed by the programme's Phase 0 through Phase 2 papers (the "internal-program" set); the remaining 95 are explicitly offered to the broader research community as the external-program research agenda.

Programme paper	Status	Catalog problems addressed	Pillar
P0-1 (this survey)	in progress	All 120 problems catalogued here; the survey itself resolves O3, O4, O7, O8, O118 (the bibliometric, taxonomic, and dependency-graph problems) through its own structure.	all
P0-2 CRB–Rate Letter	published (AIR J. Eng. Technol., 2026)	O21 partial (Kalman-ISAC single-epoch specialization); contributes foundational result cited by O10, O11, O38. First formal instance of the GNSS–ISAC isomorphism. [52]	I

Programme paper	Status	Catalog problems addressed	Pillar
P0-3 Position Paper	published (AIR J. Eng. Technol., 2026)	O1 (successor-theorem argument), O9 (historical-gap framework) — both in short form. [19]	all (framing)
P1-1 SISO Fisher–Shannon	planned	O10 (full FS boundary for SISO Gaussian specialization), O11 (FS-1 proof, SISO case), O12 (FS-2 proof, SISO case).	I
P1-2 MIMO Fisher–Shannon	planned	O10 (full MIMO boundary), O11 (FS-1 MIMO), O12 (FS-2 MIMO), O17 (extended-target MIMO).	I
P1-3 Fading-Channel ISAC	planned	O26 (fading-channel ISAC capacity across models).	I
P1-4 Kalman-ISAC Theory	planned	O21 (Kalman-ISAC steady-state, full-generality), O22 (transient response), O99 (cross-pillar Kalman-FS trilogy).	I + Synthesis
P1-5 LEO-PNT Info. Theory	planned	O50 (multi-constellation SAGIN CRB–Rate for LEO case), O52 (LEO-MEO integration), O57 (SoOp characterization).	III
P1-6 Zak-OTFS ISAC	planned	O38 (Zak-OTFS CRB–Rate region), O103 (capacity-achieving pilot density), O104 (Frontier 1 hardware).	F1 (I + II)
P2-1 HMIMO ISAC Capacity	planned	O84 (mutual coupling), O90 (EIT-compliant ISAC), O93 (RIS under EIT).	V + F5
P2-2 Near-field Positioning	planned	O87 (near-field PCV-CRB), O88 (HMIMO calibration), O92 (near-field beam-focusing CRB).	V
P2-3 AI-native ISAC Sample Complexity	planned	O69 (PAC-learning ISAC neural), O19 (Pillar I sample complexity), O100 (cross-pillar neural-FS gap).	IV + Synthesis
P2-4 Provable-Safety AI Receivers	planned	O71 (provable-safety analog), O75 (safe-fallback architecture), O111 (semantic-positioning safety).	IV + F3
P2-5 Quantum PNT	planned	O112 (quantum PNT CRB), O113 (quantum-ICL for PNT), O114 (quantum-classical hybrid SAGIN).	F4
P3-series (program continuations)	future	O14, O15 (multi-user ISAC grand challenges); O32 (native DD capacity); O50 full-generality; O68 (neural-codebook capacity); O84 (EIT mutual coupling full generality); O96, O97, O100 (synthesis grand challenges).	multiple

The table is intended as a working document for program management: it tracks which problems are actively being attacked by the program's own papers and, conversely, which problems remain open for the broader research community to address.

11. Research Roadmap and Conclusions

The preceding ten sections have argued that beyond-6G wireless has a theorem-shaped hole, that the geodetic-estimation tradition supplies a substantial fraction of the mathematics needed to fill it, and that the 120 open problems of the master catalog constitute a well-posed, presently unassigned research agenda. This closing section converts that agenda into a decade roadmap, states the dependencies that order it, restates the timing argument that makes it urgent, and issues the call to the two research communities whose engagement it requires.

11.1 The decade roadmap

The roadmap proceeds in five phases. Phase 0 (flag-planting, 2026) establishes the problem taxonomy: the present survey, the companion CRB–Rate letter [52], and the position paper [19] together define the vocabulary, prove the first operational instance of the GNSS–ISAC isomorphism, and publish the open-problem catalog. Phase 1 (foundation theorems, 2026–2028) proves the base-case results on which everything else rests: the formal construction of the Fisher–Shannon metric space $\mathscr{F}$_FS with proofs or refutations of Conjectures FS-1 and FS-2 at the matrix-valued level, the MIMO generalization, the delay-Doppler capacity theorem under the sparse-scatterer model, the Kalman-ISAC steady-state characterization of Conjecture FS-4, the LEO dual-function capacity–positioning region, and the Zak-OTFS ISAC convergence result that resolves problem O38. Phase 2 (assumption relaxation, 2027–2029) systematically removes the idealizations of the Phase 1 theorems—fading in place of AWGN, multi-user in place of point-to-point, extended in place of point targets, imperfect in place of perfect channel state information—with each relaxation drawn from the stated-limitations lists that the Phase 1 papers are required to publish. Phase 3 (cross-pillar convergence, 2028–2031) attacks the five frontiers of §9 and the synthesis problems O96 through O102, which pay off only after the single-pillar foundations exist. Phase 4 (institutionalization, 2029 onward) transitions the program from individual results to community infrastructure: dedicated workshops, graduate curricula that cross-train in geodesy and information theory, and standards contributions.

The dependency structure is explicit in the cross-reference table of §10.4. The Fisher–Shannon construction of problems O10 through O12 must precede every pillar-specific specialization; the Kalman temporal completion of O21 must precede the SAGIN steady-state problems of Pillar III and the provable-guarantee problems of Pillar IV; the delay-Doppler capacity base cases of O32 and O38 must precede the OTFS-ISAC convergence results of Frontier 1; and the synthesis problems of §8.4 presuppose all five pillars. Of the 120 catalog problems, approximately 25 are claimed by the program's own Phase 0 through Phase 2 papers; the remaining 95 are offered to the broader community, and the classification metadata of §10.1—difficulty, prerequisites, effort, spiral potential—is intended to let a doctoral student, an established group, or a funding agency select an entry point at the appropriate scale.

11.2 The window

The timing argument developed in §1 bears restating as the roadmap's governing constraint. Three institutional calendars bound the period in which foundational theory can still shape architecture: the ITU-R IMT-2030 evaluation phase runs through approximately 2029 with specifications expected in 2030; World Radiocommunication Conference 2027 will issue the spectrum allocations that determine what physics 7G can build on; and 3GPP Release 21, expected to contain the first normative 6G specification, freezes around 2028. A theorem published before these milestones can shape what gets standardized; a theorem published after them comments on what others have already built. The historical lead times of Telatar's multi-antenna capacity before 4G MIMO and Arıkan's polarization before 5G control coding both exceeded five years, which is precisely the interval separating the present survey from the 2030 horizon. The window is not closing at some indefinite future date; it is the 2026–2030 period in which this survey is published.

11.3 Call to the community and conclusions

The call is symmetric. To the geodetic and GNSS community: the beyond-6G literature is where the estimation-theoretic tools of the past half-century—the Cramér–Rao bound, dilution of precision, Kalman filtering, ambiguity-function design, multi-constellation fusion, antenna calibration—are being rediscovered under different names and, in places, with weaker guarantees; the venues of that literature are open to contributions that bring geodetic rigor to the open problems catalogued here. To the wireless-communications community: the geodesy literature contains deployed, operationally validated solutions to problems that appear in beyond-6G papers as open questions, and the investment required to access it is a short reading list rather than a second doctorate. The barrier between the two fields is bibliographic and sociological, not intellectual.

This survey has organized the foundational theory of seventh-generation wireless around five pillars, formalized the Fisher–Shannon metric space as the unifying object that appears in each, demonstrated the GNSS–ISAC isomorphism that makes six decades of geodetic estimation theory directly applicable, and published a classified catalog of 120 open problems together with the decade roadmap that orders them. None of the pillars is complete; that is the point. Every prior wireless generation was preceded by a foundational theorem, and the theorems that will play that role for 7G have not yet been proven. The argument of this survey is that they are provable, that the mathematical language in which to prove them already exists, and that the period in which proving them can still shape what gets built is now. The next Shannon may come from an unexpected field; this survey is an invitation to help write the result that earns the comparison.

12. Limitations

A survey that argues a thesis and publishes a research agenda incurs two distinct classes of limitation: those affecting the argument, and those affecting the agenda. This section states both. Several are noted in passing where they arise — the qualifications of §2.5, the scope statement of §3.3, the provenance note of §10.1 — and are collected here so that a reader can weigh them together rather than encountering them piecemeal.

12.1 The central thesis is contestable

The thesis of §2.5 is that beyond-6G wireless has entered normative standardization without a foundational theorem of the kind that preceded earlier generations. Two objections deserve statement in their strongest form. The first is that the historical pattern may be an artifact of retrospective selection: Shannon, Telatar and Arıkan are identifiable as foundational in hindsight, and it is possible that the results which will later be identified as foundational for 6G already exist in the literature surveyed here — the Xiong-Liu CRB–Rate characterization [9] being the most plausible candidate — and are simply not yet recognized as such. On this reading the gap this survey identifies is a gap in recognition rather than in mathematics. The qualification of §1.1 narrows the pattern to generations that changed what the physical layer optimizes, which is the author's answer to this objection, but the answer is an argument rather than a demonstration. The second objection is that standardization may simply not require foundational theory: 3GPP has repeatedly specified systems in advance of the analysis that later explained them, and a reader who regards that as the normal and sufficient order of operations will regard the urgency argument of §1.3 and §11.2 as overstated. The survey does not refute this position. It asserts, without proving, that specifications resting on empirical rather than provable performance claims are harder to defend, extend and port — a claim about engineering practice for which no evidence is offered here.

12.2 The catalog metadata are judgments, not measurements

The effort estimates and spiral-potential figures of Table 10.1 are unvalidated single-assessor judgments, as stated in §10.1. Two hundred and forty numbers are presented in tabular form, which confers an appearance of precision that their derivation does not support. No second assessor has scored them and no calibration has been performed. The difficulty classification is better founded, since §10.1 gives operational criteria, but it too has been applied by one person without inter-rater checking, and the boundary between Hard and Grand Challenge is in several cases a judgment that another researcher could reasonably draw elsewhere. Readers using the catalog for portfolio selection should treat the classification as a first-order sort rather than as a ranking, and should expect the effort figures to be wrong by a factor of two in individual cases even if the aggregate is approximately right.

12.3 The pillar partition is a choice, and the coverage is bounded

The five-pillar structure follows the selection principle stated in §1.4 and is not a partition of the beyond-6G research agenda. Terahertz propagation, physical-layer security, energy-efficiency theory and semantic communication are excluded as pillars, and a survey organized on a different principle would produce a different and equally defensible taxonomy. Two consequences follow. Problems that live wholly inside an excluded area are absent from the catalog, so the 120 problems registered here should not be read as an enumeration of what remains open in beyond-6G theory. And the cross-pillar frontiers of §9 are frontiers relative to this partition; a different partition would locate them differently. The literature reviewed is also bounded in ways worth recording: it is English-language throughout; it includes eight arXiv preprints whose results have not completed peer review; and its coverage of the geodetic literature is weighted toward the GNSS and reference-frame material closest to the author's own training, which means that adjacent estimation-theoretic traditions — inertial navigation, photogrammetry, and the radar-tracking literature — are represented more thinly than their relevance to Pillars I and III would justify.

12.4 The framework is proposed rather than proved

Two assumptions underlie the survey and are stated here as assumptions. The first is that the geodetic estimation toolkit transfers to the beyond-6G setting: this premise is argued in §1.2, §3.3 and §8.3 and proved in one instance [52], but the pillar-by-pillar bridges of §§4–7 assert the correspondence rather than deriving it, and its generality is assumed throughout. The second is that the Fisher–Shannon metric space introduced in §3.4 is a proposed organizing object, not an established one. All five of its substantive properties are stated as conjectures FS-1 through FS-5, and none is proved in this survey. If FS-1 or FS-2 fails — if Gaussian inputs do not achieve the boundary, or if the region is not convex under the cone-product order — the unification of §8.1 would require reformulation, though the individual pillar problems would survive. The single result on which the strongest claim of this survey depends, the closed-form SISO bridge of [52], is likewise the author's own, and its generalization is the subject of the open problems rather than of any result presented here.

13. Acknowledgments

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

References

[1] C. E. Shannon, "A Mathematical Theory of Communication," Bell Syst. Tech. J., vol. 27, no. 3, pp. 379–423, Jul. 1948, doi: 10.1002/j.1538-7305.1948.tb01338.x.

[2] E. Telatar, "Capacity of Multi-antenna Gaussian channels," Eur. Trans. Telecommun., vol. 10, no. 6, pp. 585–595, Nov.–Dec. 1999 (online 12 Sep. 2008), doi: 10.1002/ett.4460100604.

[3] E. Arikan, "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, doi: 10.1109/TIT.2009.2021379.

[4] 3GPP, "NR; Multiplexing and channel coding (polar codes in Release 15 control channels)," 3rd Generation Partnership Project, Tech. Spec. TS 38.212, Rel-15, Jun. 2018. [Online]. Available: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3214 (Accessed: Aug. 22, 2026).

[5] ITU-R, "Framework and overall objectives of the future development of IMT for 2030 and beyond," International Telecommunication Union, Recommendation ITU-R M.2160-0, Nov. 2023. [Online]. Available: https://www.itu.int/dms_pubrec/itu-r/rec/m/R-REC-M.2160-0-202311-I!!PDF-E.pdf (Accessed: Aug. 22, 2026).

[6] ITU, "Final Acts: World Radiocommunication Conference (WRC-23)," International Telecommunication Union, Geneva, Switzerland, Dec. 2023. [Online]. Available: http://handle.itu.int/11.1002/pub/8225d4fb-en (Accessed: Aug. 22, 2026).

[7] 3GPP, "Release 21 timeline agreed at TSGs#112," 3rd Generation Partnership Project, RP-260868 / SP-260595 / CP-261259, Jun. 2026. [Online]. Available: https://www.3gpp.org/news-events/3gpp-news/rel21-timeline (Accessed: Aug. 22, 2026).

[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., doi: 10.1109/JSAC.2022.3156632.

[9] 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, Sept. 2023, doi: 10.1109/TIT.2023.3284449.

[10] J. An, C. Yuen, C. Huang, M. Debbah, H. V. Poor, and L. Hanzo, "A Tutorial on Holographic MIMO Communications—Part I: Channel Modeling and Channel Estimation," IEEE Commun. Lett., vol. 27, no. 7, pp. 1664–1668, Jul. 2023, doi: 10.1109/LCOMM.2023.3278683.

[11] C. -X. Wang, Y. Yang, J. Huang, X. Gao, T. J. Cui and L. Hanzo, "Electromagnetic Information Theory: Fundamentals and Applications for 6G Wireless Communication Systems," IEEE Wireless Commun., vol. 31, no. 5, pp. 279–286, Oct. 2024, doi: 10.1109/MWC.013.2300399.

[12] E. D. Kaplan and C. J. Hegarty, Eds., Understanding GPS/GNSS: Principles and Applications, 3rd ed. Norwood, MA, USA: Artech House, 2017.

[13] P. Misra and P. Enge, Global Positioning System: Signals, Measurements, and Performance, 2nd ed. Lincoln, MA, USA: Ganga-Jamuna Press, 2011.

[14] P. J. G. Teunissen and A. Kleusberg, Eds., GPS for Geodesy, 2nd ed. Berlin, Germany: Springer, 1998.

[15] 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).

[16] 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.

[17] 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.

[18] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Physical channels and modulation (MIMO in Release 8)," 3rd Generation Partnership Project, Tech. Spec. TS 36.211, Rel-8, Dec. 2008. [Online]. Available:

https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2425 (Accessed: Aug. 22, 2026).

[19] M. Hawarey, "Why 7G Needs New Theorems: A Geodesy-based Perspective on Beyond-6G Foundations," AIR J. Eng. Technol., 2026, Art. no. AIRJET2026811, doi: 10.65737/AIRJET2026811.

[20] ITU-R, "Technical performance requirements for IMT-2030 radio interface(s)," International Telecommunication Union, draft new Report, Working Party 5D, Document 5/116, completed Feb. 2026, submitted to Study Group 5 for approval Dec. 2026. [Online]. Available: https://www.itu.int/md/R23-SG05-C-0116/en (Accessed: Aug. 23, 2026).

[21] 3GPP, "Study on artificial intelligence (AI) / machine learning (ML) for NR air interface," 3rd Generation Partnership Project, Tech. Rep. TR 38.843, Rel-18, v18.0.0, Dec. 2023. [Online]. Available: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3983 (Accessed: Aug. 23, 2026).

[22] 3GPP, "Release 20: final 5G-Advanced release and 6G study items," 3rd Generation Partnership Project, release description and Work Plan SP-260360, Mar. 2026. [Online]. Available: https://www.3glteinfo.com/releases/release-20/ (Accessed: Aug. 23, 2026).

[23] 3GPP, "Study on integrated sensing and communication," 3rd Generation Partnership Project, Tech. Rep. TR 22.837, v19.4.0, Rel-19, 2024. [Online]. Available: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=4044 (Accessed: Aug. 23, 2026).

[24] M. Alsabah, M. A. Naser, B. M. Mahmmod, S. H. Abdulhussain, M. R. Eissa, A. Al-Baidhani, N. K. Noordin, S. M. Sait, K. A. Al-Utaibi, and F. Hashim, "6G Wireless Communications Networks: A Comprehensive Survey," IEEE Access, vol. 9, pp. 148191–148243, 2021, doi: 10.1109/ACCESS.2021.3124812.10.1109/ACCESS.2021.3124812

[25] M. Giordani, M. Polese, M. Mezzavilla, S. Rangan, and M. Zorzi, "Toward 6G Networks: Use Cases and Technologies," IEEE Commun. Mag., vol. 58, no. 3, pp. 55–61, Mar. 2020, doi: 10.1109/MCOM.001.1900411.

[26] M. A. Uusitalo, P. Rugeland, M. R. Boldi, E. C. Strinati, P. Demestichas, M. Ericson, G. P. Fettweis, M. C. Filippou, A. Gati, M-H. Hamon, M. Hoffmann, M. Latva-Aho, A. Parssinen, B. Richerzhagen, H. Schotten, T. Svensson, G. Wikstrom, H. Wymeersch, V. Ziegler, and Y. Zou, "6G Vision, Value, Use Cases and Technologies From European 6G Flagship Project Hexa-X," IEEE Access, vol. 9, pp. 160004–160020, 2021, doi: 10.1109/ACCESS.2021.3130030.

[27] D. C. Nguyen, M. Ding, P. N. Pathirana, A. Seneviratne, J. Li, and D. Niyato, "6G Internet of Things: A Comprehensive Survey," IEEE Internet Things J., vol. 9, no. 1, pp. 359–383, Jan. 2022, doi: 10.1109/JIOT.2021.3103320.

[28] 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, Apr. 2025, doi: 10.1109/OJCOMS.2025.3560035.

[29] E. Hossain and A. Vera-Rivera, "6G Cellular Networks: Mapping the Landscape for the IMT-2030 Framework," arXiv preprint arXiv:2501.14552, Jan. 2025. [Online]. Available: https://arxiv.org/abs/2501.14552 (Accessed: Aug. 23, 2026).

[30] J. Zhang, W. Lu, C. Xing, N. Zhao, N. Al-Dhahir, G. K. Karagiannidis, and X. Yang, "Intelligent Integrated Sensing and Communication: A Survey," Sci. China Inf. Sci., vol. 68, no. 3, Art. no. 131301, Mar. 2025, doi: 10.1007/s11432-024-4205-8.

[31] A. Liu, Z. Huang, M. Li, Y. Wan, W. Li, T. X. Han, C. Liu, R. Du, D. K. P. Tan, J. Lu, Y. Shen, F. Colone, and K. Chetty, "A Survey on Fundamental Limits of Integrated Sensing and Communication," IEEE Commun. Surveys Tuts., vol. 24, no. 2, pp. 994–1034, Feb. 2022, doi: 10.1109/COMST.2022.3149272.

[32] 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, doi: 10.1109/WCNC.2017.7925924.

[33] M. Ubadah, S. K. Mohammed, R. Hadani, S. Kons, A. Chockalingam, and R. Calderbank, "Zak-OTFS to Integrate Sensing the I/O Relation and Data Communication," arXiv preprint arXiv:2404.04182, 2024. [Online]. Available: https://arxiv.org/abs/2404.04182 (Accessed: Aug. 23, 2026).

[34] J. Liu, Y. Shi, Z. M. Fadlullah, and N. Kato, "Space-Air-Ground Integrated Network: A Survey," IEEE Commun. Surveys Tuts., vol. 20, no. 4, pp. 2714–2741, May 2018, doi: 10.1109/COMST.2018.2841996.

[35] N. Cheng, J. He, Z. Yin, C. Zhou, H. Wu, F. Lyu, H. Zhou, X. Shen, "6G service-oriented space-air-ground integrated network: A survey," Chin. J. Aeronaut., vol. 35, no. 9, pp. 1–18, Sep. 2022, doi: 10.1016/j.cja.2021.12.013.

[36] Q. Chen, Z. Guo, W. Meng, S. Han, C. Li, and T. Q. S. Quek, "A Survey on Resource Management in Joint Communication and Computing-Embedded SAGIN," arXiv preprint arXiv:2403.17400, 2024. [Online]. Available: https://arxiv.org/abs/2403.17400 (Accessed: Aug. 23, 2026).

[37] F. Sun, A. T. Abebe, H. Ji, and E. Onggosanusi, "6G AI/ML for Physical Layer: Part II – CSI Compression with JSCM," (technical blog and research program), Samsung Research, 2025. [Online]. Available: https://research.samsung.com/blog/6G-AI-ML-for-Physical-layer-Part-II-CSI-Compression-with-JSCM (Accessed: Aug. 23, 2026).

[38] Q. Cui, X. You, N. Wei, G. Nan, X. Zhang, J. Zhang, X. Lyu, M. Ai, X. Tao, Z. Feng, P. Zhang, Q. Wu, M. Tao, Y. Huang, C. Huang, G. Liu, C. Peng, Z. Pan, T. Sun, D. Niyato, T. Chen, M. K. Khan, A. Jamalipour, M. Guizani, and C. Yuen, "Overview of AI and communication for 6G network: fundamentals, challenges, and future research opportunities," Sci. China Inf. Sci., vol. 68, no. 7, Art. no. 171301, Jul. 2025, doi: 10.1007/s11432-024-4337-1.

[39] F. C. Ogenyi, C. N. Ugwu, and O. P.-C. Ugwu, "A comprehensive review of AI-native 6G: integrating semantic communications, reconfigurable intelligent surfaces, and edge intelligence for

next-generation connectivity," Front. Commun. Netw., Sep. 2025, 6:1655410. doi: 10.3389/frcmn.2025.1655410.

[40] M. D. Migliore, "On Electromagnetics and Information Theory," IEEE Trans. Antennas Propag., vol. 56, no. 10, pp. 3188–3200, Oct. 2008, doi: 10.1109/TAP.2008.929444.10.1109/TAP.2008.929444

[41] A. Pizzo, L. Sanguinetti, and T. L. Marzetta, "Fourier Plane-Wave Series Expansion for Holographic MIMO Communications," IEEE Trans. Wireless Commun., vol. 21, no. 9, pp. 6890–6905, Mar. 2022, doi: 10.1109/TWC.2022.3152965.10.1109/TWC.2022.3152965

[42] J. An, C. Yuen, C. Huang, M. Debbah, H. V. Poor and L. Hanzo, "A Tutorial on Holographic MIMO Communications—Part II: Performance Analysis and Holographic Beamforming," IEEE Commun. Lett., vol. 27, no. 7, pp. 1669–1673, Jul. 2023, doi: 10.1109/LCOMM.2023.3278682.

[43] M. Kobayashi, G. Caire, and G. Kramer, "Joint state sensing and communication: Optimal tradeoff for a memoryless case," in Proc. IEEE Int. Symp. Inf. Theory (ISIT), Vail, CO, USA, Jun. 2018, doi: 10.1109/ISIT.2018.8437621.

[44] M. Ahmadipour, M. Kobayashi, M. Wigger, and G. Caire, "An information-theoretic approach to joint sensing and communication," IEEE Trans. Inf. Theory, vol. 70, no. 2, pp. 1124–1146, Feb. 2024, doi: 10.1109/TIT.2022.3176139.

[45] F. Lampel, H. Joudeh, A. Alvarado, and F. M. J. Willems, "Orthogonal time frequency space modulation based on the discrete Zak transform," Entropy, vol. 24, no. 12, Art. no. 1704, Nov. 2022, doi: 10.3390/e24121704.

[46] F. Liu, Y. Xiong, K. Wan, T. X. Han, and G. Caire, "Deterministic-random tradeoff of integrated sensing and communications in Gaussian channels: A rate-distortion perspective," arXiv preprint arXiv:2212.10897, 2022. [Online]. Available: https://arxiv.org/abs/2212.10897 (Accessed: Aug. 23, 2026).

[47] 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, Mar. 2024, doi: 10.1109/MBITS.2024.3376638.

[48] X. Li, Z. Tang, V. C. Andrei, U. J. Monich, F. Liu, and H. Boche, "Computation of Capacity-Distortion-Cost Functions for Continuous Memoryless Channels," arXiv preprint arXiv:2504.20285, Apr. 2025. [Online]. Available: https://arxiv.org/abs/2504.20285 (Accessed: Aug. 23, 2026).

[49] K. Dong, F. Liu, Y. Xiong, Y. Cui, and Y. C. Eldar, "Fundamental Limits of Communication-Assisted Sensing in ISAC Systems," arXiv preprint arXiv:2404.08188. [Online]. Available: https://ar5iv.labs.arxiv.org/html/2404.08188 (Accessed: Aug. 23, 2026).

[50] G. Peng and Y. Wu, "Joint source-channel coding for ISAC: Distortion tradeoffs and separation theorems," arXiv preprint arXiv:2601.10470, Jan. 2026. [Online]. Available: https://arxiv.org/abs/2601.10470 (Accessed: Aug. 23, 2026).

[51] C. Ouyang, Y. Liu, H. Yang, and N. Al-Dhahir, "Integrated Sensing and Communications: A Mutual Information-Based Framework," IEEE Commun. Mag., vol. 61, no. 5, pp. 26–32, May 2023, doi: 10.1109/MCOM.001.2200493.

[52] M. Hawarey, "On the fundamental CRB–Rate tradeoff for GNSS-integrated sensing and communication," AIR J. Eng. Technol., Jun. 2026, Art. no. AIRJET2026793, doi: 10.65737/AIRJET2026793.

[53] P. M. Woodward, Probability and Information Theory, with Applications to Radar. London, U.K.: Pergamon Press, 1953.

[54] S. K. Mohammed, R. Hadani, A. Chockalingam, and R. Calderbank, "OTFS—Predictability in the Delay-Doppler Domain and Its Value to Communication and Radar Sensing," IEEE BITS the Information Theory Magazine, vol. 3, no. 2, pp. 7–31, Jun. 2023, doi: 10.1109/MBITS.2023.3319595.

[55] N. Mehrotra, S. R. Mattu, and R. Calderbank, "Zak-OTFS with spread carrier waveforms," IEEE Wireless Commun. Lett., vol. 14, no. 10, pp. 3244–3248, Oct. 2025, doi: 10.1109/LWC.2025.3590254.

[56] J. Jayachandran, I. A. Khan, S. K. Mohammed, R. Hadani, A. Chockalingam, and R. Calderbank, "Zak-OTFS with Interleaved Pilots to Extend the Region of Predictable Operation," IEEE Trans. Veh. Technol., vol. 74, no. 11, pp. 17545–17559, Jun. 2025, doi: 10.1109/TVT.2025.3579394.

[57] IEEE Communications Society, "Best Readings in Orthogonal Time Frequency Space (OTFS) and Delay Doppler Signal Processing," IEEE ComSoc Educational Program, 2022–2025. [Online]. Available: https://www.comsoc.org/publications/best-readings/orthogonal-time-frequency-space-otfs-and-delay-doppler-signal-processing (Accessed: Aug. 23, 2026).

[58] A. J. E. M. Janssen, "The Zak transform: A signal transform for sampled time-continuous signals," Philips J. Res., vol. 43, no. 1, pp. 23–69, 1988.

[59] A. Bakambekova, N. Kouzayha, and T. Al-Naffouri, "On the Interplay of Artificial Intelligence and Space-Air-Ground Integrated Networks: A Survey," IEEE Open J. Commun. Soc., vol. 5, pp. 4613–4673, Jul. 2024, doi: 10.1109/OJCOMS.2024.3429198.10.1109/OJCOMS.2024.342919810.1109/OJCOMS.2024.3429198

[60] L. Toka, E. A. Papp, T. Cinkler, I. Gódor, and L. Hévizi, "Dimensioning space-air-ground integrated networks for in-flight 6G slice orchestration," Veh. Commun., vol. 51, 2025, doi: 10.1016/j.vehcom.2024.100866.

[61] N. Zhang, S. Zhang, P. Yang, O. Alhussein, W. Zhuang and X. S. Shen, "Software Defined Space-Air-Ground Integrated Vehicular Networks: Challenges and Solutions," IEEE Commun. Mag., vol. 55, no. 7, pp. 101–109, Jul. 2017, doi: 10.1109/MCOM.2017.1601156.

[62] J. Chen, H. Zhang, and Z. Xie, "Space-Air-Ground Integrated Network (SAGIN): A survey," arXiv preprint arXiv:2307.14697, Jul. 2023. [Online]. Available: https://arxiv.org/abs/2307.14697 (Accessed: Aug. 23, 2026).

[63] O. Kodheli, E. Lagunas, N. Maturo, S. K. Sharma, B. Shankar, J. F. M. Montoya, J. C. M. Duncan, D. Spano, S. Chatzinotas, S. Kisseleff, J. Querol, L. Lei, T. X. Vu, and G. Goussetis, "Satellite Communications in the New Space Era: A Survey and Future Challenges," IEEE Commun. Sur. & Tuts, vol. 23, no. 1, pp. 70–109, Oct. 2020, doi: 10.1109/COMST.2020.3028247.

[64] Z. M. Kassas, N. Khairallah, and S. Kozhaya, "Ad astra: Simultaneous Tracking and Navigation with Megaconstellation LEO Satellites," IEEE Aerosp. Electron. Syst. Mag., vol. 39, no. 9, pp. 46–71, Jan. 2024, doi: 10.1109/MAES.2023.3267440.

[65] T. G. R. Reid, T. Walter, P. K. Enge, D. Lawrence, H. S. Cobb, G. Gutt, M. O'Connor, and D. Whelan, "Navigation from Low Earth Orbit — Part 1: Concept, Current Capability, and Future Promise," in Position, Navigation, and Timing Technologies in the 21st Century: Integrated Satellite Navigation, Sensor Systems, and Civil Applications Volume 2. Hoboken, NJ, USA: Wiley-IEEE Press, 2020, doi: 10.1002/9781119458555.ch43a.

[66] Z. M. Kassas, "Navigation from Low-Earth Orbit — Part 2: Models, Implementation, and Performance," in Position, Navigation, and Timing Technologies in the 21st Century: Integrated Satellite Navigation, Sensor Systems, and Civil Applications Volume 2. Hoboken, NJ, USA: Wiley-IEEE Press, 2020, doi: 10.1002/9781119458555.ch43b.

[67] European Space Agency, "Moonlight Lunar Communications and Navigation System (LCNS)," ESA Programme Description, 2020-2024. [Online]. Available: https://www.esa.int/Applications/Connectivity_and_Secure_Communications/Moonlight (Accessed: Aug. 23, 2026).

[68] M. Hawarey, "Travelling ionospheric disturbance over California mid 2000," Nonlin. Processes Geophys., vol. 13, no. 1, pp. 1–7, Jan. 2006, doi: 10.5194/npg-13-1-2006.

[69] T. O'Shea and J. Hoydis, "An Introduction to Deep Learning for the Physical Layer," IEEE Trans. Cogn. Commun. Netw., vol. 3, no. 4, pp. 563–575, Oct. 2017, doi: 10.1109/TCCN.2017.2758370.

[70] S. Dorner, S. Cammerer, J. Hoydis, and S. ten Brink, "Deep Learning Based Communication Over the Air," IEEE J. Sel. Topics Signal Process., vol. 12, no. 1, pp. 132–143, Dec. 2017, doi: 10.1109/JSTSP.2017.2784180.

[71] D. Gunduz, Z. Qin, I. E. Aguerri, H. S. Dhillon, Z. Yang, A. Yener, K. K. Wong, and C.-B. Chae, "Beyond Transmitting Bits: Context, Semantics, and Task-Oriented Communications," IEEE J. Sel. Areas Commun., vol. 41, no. 1, pp. 5–41, Nov. 2022, doi: 10.1109/JSAC.2022.3223408.

[72] NVIDIA Corporation, "Sionna: a hardware-accelerated differentiable open-source library for research on communication systems," NVIDIA Corp. Website. [Online]. Available: https://nvlabs.github.io/sionna (Accessed: Aug. 23, 2026).

[73] DeepSig Inc., "OmniPHY: Machine Learning in the Physical (PHY) Layer," DeepSig Corp. Website. [Online]. Available: https://www.deepsig.ai/omniphy-5g/ (Accessed: Aug. 23, 2026).

[74] MathWorks, "AI for 6G," MathWorks Corp. Website. [Online]. Available: https://www.mathworks.com/help/5g/ai-for-6g.html (Accessed: Aug. 23, 2026).

[75] R. Wiesmayr, C. Dick, J. Hoydis, and C. Studer, "DUIDD: Deep-Unfolded Interleaved Detection and Decoding for MIMO Wireless Systems," in Proc. 56th Asilomar Conf. Signals, Syst., Comput. (ACSSC), Pacific Grove, CA, USA, Oct. 31–Nov. 2, 2022, doi: 10.1109/IEEECONF56349.2022.10051944.

[76] M. Hawarey, "Characterizing In-Context Learning: When Can Transformers Match Standard Learning Algorithms?" AIR J. Math. Comput. Sci., Feb. 2026, Art. no. AIRMCS2026277, doi: 10.65737/AIRMCS2026277.

[77] M. Hawarey, "Agents of Context: A Methodological Critique and Counter-Evidence Analysis of Adversarial Red-Teaming Claims for Autonomous AI Agents," AIR J. Interdiscip. Res., Mar. 2026, Art. no. AIRJIR2026369, doi: 10.65737/AIRJIR2026369.

[78] M. Hawarey, "Chain-of-Thought ICL for GeoAI: Breaking the Hardness Barrier," AIR J. Math. Comput. Sci., Mar. 2026, Art. no. AIRMCS2026482, doi: 10.65737/AIRMCS2026482.

[79] M. Hawarey, "ICL Characterization of Geo-Foundation Models," AIR J. Math. Comput. Sci., Feb. 2026, Art. no. AIRMCS2026360, doi: 10.65737/AIRMCS2026360.

[80] M. Hawarey, "ICL Characterization of Climate Foundation Models: When Can Transformers Learn Weather and Climate?" AIR J. Math. Comput. Sci., Mar. 2026, Art. no. AIRMCS2026405, doi: 10.65737/AIRMCS2026405.

[81] M. Hawarey, "ICL Characterization of Multi-Modal Geo-Foundation Models: When Can Vision-Language Transformers Learn Geospatial Tasks?" AIR J. Math. Comput. Sci., Mar. 2026, Art. no. AIRMCS2026446, doi: 10.65737/AIRMCS2026446.

[82] M. Hawarey, "Quantum-Enhanced In-Context Learning for Geopotential Field Estimation: A Theoretical Framework," AIR J. Math. Comput. Sci., Feb. 2026, Art. no. AIRMCS2026322, doi: 10.65737/AIRMCS2026322.

[83] M. Franceschetti, M. D. Migliore, and P. Minero, "The Capacity of Wireless Networks: Information-Theoretic and Physical Limits," IEEE Trans. Inf. Theory, vol. 55, no. 8, pp. 3413–3424, Jul. 2009, doi: 10.1109/TIT.2009.2023705.

[84] A. Pizzo, A. d. J. Torres, L. Sanguinetti, and T. L. Marzetta, "Nyquist Sampling and Degrees of Freedom of Electromagnetic Fields," IEEE Trans. Signal Process., vol. 70, pp. 3935–3947, Jun. 2022, doi: 10.1109/TSP.2022.3183445.

[85] C. Huang, S. Hu, G. C. Alexandropoulos, A. Zappone, C. Yuen, R. Zhang, M. Di Renzo, and M. Debbah, "Holographic MIMO Surfaces for 6G Wireless Networks: Opportunities, Challenges, and Trends," IEEE Wireless Commun., vol. 27, no. 5, pp. 118–125, Oct. 2020, doi: 10.1109/MWC.001.1900534.

[86] D. Dardari and N. Decarli, "Holographic Communication Using Intelligent Surfaces," IEEE Commun. Mag., vol. 59, no. 6, pp. 35–41, Jun. 2021, doi: 10.1109/MCOM.001.2001156.

[87] T. Gong, L. Wei, C. Huang, Z. Yang, J. He, and M. Debbah, "Holographic MIMO Communications With Arbitrary Surface Placements: Near-Field LoS Channel Model and Capacity Limit," IEEE J. Sel. Areas Commun., vol. 42, no. 6, pp. 1549–1566, Apr. 2024, doi: 10.1109/JSAC.2024.3389126.

[88] J. Zhu, V. Y. F. Tan, and L. Dai, "MIMO Capacity Analysis and Channel Estimation for Electromagnetic Information Theory," IEEE J. Sel. Areas Commun., vol. 44, pp. 3302–3316, Jan. 2026, doi: 10.1109/JSAC.2026.3654525.

[89] X. Q. Chen, L. Zhang, Y. N. Zheng, S. Liu, Z. R. Huang, J. C. Liang, M. D. Renzo, V. Galdi, and T. J. Cui, "Integrated sensing and communication based on space-time-coding metasurfaces," Nat. Commun., vol. 16, Art. no. 1836, Feb. 2025, doi: 10.1038/s41467-025-57137-6.

[90] Y. Liu, C. Ouyang, Z. Wang, J. Xu, X. Mu and A. L. Swindlehurst, "Near-Field Communications: A Comprehensive Survey," IEEE Commun. Surveys Tuts., vol. 27, no. 3, pp. 1687–1728, Jun. 2025, doi: 10.1109/COMST.2024.3475884.

[91] IEEE Information Theory Society, "Call for papers: Electromagnetic information theory," special issue, IEEE Trans. Inf. Theory / IEEE J. Sel. Areas Inf. Theory, 2024. [Online]. Available: https://www.itsoc.org (Accessed: Aug. 22, 2026).

[92] 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, pp. 781–798, Apr. 2007, doi: 10.1007/s00190-007-0148-y.