

Quantum Sensing as an Optimization–Inference Stack

Sridhar Tayur

Quantum Technologies Group, Carnegie Mellon University, Pittsburgh, PA 15213, USA

June 2026

Revised August 2026

Revised October 2026

Abstract

Every data science pipeline begins with a measurement. When that measurement is made by a quantum mechanical system—exploiting superposition, entanglement, and coherence—the sensor can in principle surpass the classical shot-noise precision limit. Atomic clocks already underpin GPS; magnetometers inform epilepsy-surgery planning; single-photon LiDAR systems detect methane leaks at parts-per-million-meter (ppm·m) path-integrated concentration from ranges up to about 200 meters. This article has two purposes. The first is to serve as a primer for the INFORMS data science community: the author argues that the two central design problems in quantum sensing—*what state should the probe be in?* and *what measurement should be performed?*—can often be formulated or bounded using semidefinite programs (SDPs), develops both design problems from first principles, and shows how the resulting bounds constrain every other quantity in the sensing pipeline. The second purpose is to map a research agenda connecting the quantum layer to familiar data science tools: adaptive bandits, reinforcement learning, inverse problems, stochastic programming, and ML surrogates. An end-to-end methane emissions case study traces the complete stack from photon physics through SDP design to operational crew dispatch, providing the paper’s running quantum-versus-classical benchmark.

Keywords: semidefinite programming; quantum Fisher information; Bayesian estimation; bandits; reinforcement learning; adaptive measurement; quantum LiDAR; methane detection

1 Introduction

In a hospital in London, a patient with drug-resistant epilepsy wears a lightweight helmet lined with optically pumped magnetometers (OPMs); as neurons fire, the helmet records the resulting femtotesla magnetic fields and a reconstruction algorithm locates the seizure focus to within a few millimeters—without anesthesia, without a rigid scanner, and without requiring the patient to hold still [Boto et al., 2018]. Thousands of miles away, a drone carrying a single-photon LiDAR system surveys a natural gas pipeline, measuring methane plumes at parts-per-million-meter (ppm·m) path-integrated concentration from ranges up to about 200 meters and streaming a three-dimensional map of emission rates to a cloud processing pipeline that, according to company reports, dispatches repair crews within hours [MIT News Office, 2025, QLM Technology, 2025]. Both systems operate today. Both generate data—neuronal current maps, volumetric gas concentrations—that is classical in form but quantum in origin, and both require estimation, inversion, and optimization algorithms to convert raw photon counts or voltage signals into actionable decisions.

The *optimization–inference stack*, shown schematically in Figure 1, organizes the journey from quantum hardware to operational decision into three layers. The **quantum layer** encompasses everything the experimenter controls before a classical number is recorded: the choice of probe

state ρ_0 , its interaction with the signal Hamiltonian $H_s = \phi G$, the decoherence it suffers during that interaction, and the measurement applied at the end. The **bounds layer** characterizes—before any experiment is conducted—the best probe state and measurement that physics permits for a given noise model, often via semidefinite or related conic-program formulations. The resulting quantum Cramér–Rao bound (QCRB) is a lower bound on local unbiased-estimator variance—equivalently, a best-case precision benchmark—that every downstream algorithm must respect. The **classical layer** is where operations research and management science (OR/MS) methods already have a home: raw measurement outcomes become parameter estimates, uncertainty statements, and ultimately decisions about where to send a repair crew or where to place a surgical incision.

A key message is that the semidefinite programs (SDP) and non-SDP problems cataloged here are not methodologically novel to the INFORMS data science community. Multi-armed bandits with endogenous learning curves have been used for split liver transplantation allocation [Tang et al., 2025a]; active sequential hypothesis testing under assay budgets has been developed for liquid biopsy diagnostics [Gan et al., 2025]; and analytical decision support under irreversible resource constraints has been applied to organ allocation [Tang et al., 2025b]. Reinforcement learning (RL) is likewise a well-established sequential-decision-making framework within operations research [Powell, 2022], and simulation-based inference and neural posterior estimation are increasingly standard machine learning (ML) tools wherever a forward model is cheap to simulate but the likelihood is intractable [Cranmer et al., 2020]. What quantum sensing adds is unusually clean problem anatomy: exact forward models given by the Lindblad equation, SDP-certified precision ceilings, and a direct information-theoretic correspondence between algorithmic reward and physical information gain. Accordingly, the paper is best read not as a technical contribution to SDP methodology, but as a position paper that frames quantum sensing as a rich and largely untapped source of optimization and inference problems for the INFORMS data science community.

Scope and Organization

Section 2 develops the theoretical framework and two sensing-design optimizations, with SDP or conic-program formulations or bounds in important model classes, with an NV-center magnetometer as the running example threaded through each subsection. Section 3 surveys hardware by sensing modality. Section 4 treats sensing algorithms, including formal bandit, RL, and inverse-problem formulations. Section 5 covers applications. Section 6 provides an end-to-end case study of methane emissions monitoring, tracing the full stack from photon physics through SDP design, ML reconstruction, and operational dispatch—including an explicit classical-vs-quantum comparison. Section 7 covers benchmarking. Section 8 maps research opportunities for the INFORMS data science community. Section 9 concludes. Appendix A collects notation, Appendix B gives the full Ramsey derivation, SQL/HL limits with proofs, and the Cramér–Rao gap decomposition, Appendix C surveys the global hardware landscape, and Appendix D provides a glossary of quantum-sensing acronyms.

Table 1 maps each quantum sensing topic to its mathematical toolkit and key references.

To avoid conflating aspirations with achievements, each hardware platform and application is tagged with one of three maturity levels:

- **[Deployed]** — in routine operational use at scale over multiple years.
- **[Field trial]** — demonstrated outside the laboratory with independent evaluation or regulatory engagement, not yet in routine deployment.
- **[Lab/prototype]** — principle demonstrated in controlled laboratory conditions; engineering




Figure 1: The quantum sensing pipeline as an optimization–inference stack. **Quantum layer** (top): the probe state ρ_0 is prepared, exposed to the signal Hamiltonian $H_s = \phi G$, undergoes noisy evolution $\mathcal{E}_\phi(\rho_0)$ (signal encoding plus the noise map $\mathcal{E}$; Section 2.7), and is measured by a positive operator-valued measure (POVM) $\{E_i\}$ —the general description of a quantum measurement, with outcome probabilities $p_i = \text{Tr}(\rho E_i)$; see the POVM discussion in Section 2.2. **Bounds layer** (center): probe design (Optimization 1) chooses ρ_0 to maximize the quantum Fisher information F_Q^* ; measurement design (Optimization 2) uses the resulting ρ_0^* to choose a measurement POVM that minimizes the weighted estimation error $\text{Tr}(W\Sigma)$, subject to multiparameter incompatibility. Together they establish the QCRB, which constrains the achievable estimation variance; this variance constraint propagates through uncertainty quantification and thereby informs downstream decisions. **Classical layer** (bottom): raw measurement outcomes flow to parameter estimation, uncertainty quantification, and downstream operational decisions. Solid arrows carry data; dashed arrows carry bounds and design prescriptions from the bounds layer.

path to field use not yet established.

2 Theoretical Foundations

2.1 Quantum States, Operators, and Measurement

A quantum system is described by a state vector $|\psi\rangle$ in a complex Hilbert space $\mathcal{H}$. For a *qubit*—the two-level quantum analogue of a classical bit— $\mathcal{H} = \mathbb{C}^2$ and any pure state is

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \quad |\alpha|^2 + |\beta|^2 = 1.$$

Here $|0\rangle$ and $|1\rangle$ are the two basis states (analogous to 0 and 1 in a classical bit), and the coefficients $\alpha, \beta \in \mathbb{C}$ are probability amplitudes: the sensor is in *both* states simultaneously until a measurement is made. A *density operator* ρ generalizes this to statistical mixtures: $\rho \succeq 0$, $\text{Tr}(\rho) = 1$. Here $\text{Tr}(\cdot)$ denotes the matrix trace (sum of diagonal entries), and the positive-semidefinite (PSD) cone is exactly the setting of semidefinite programming—density operators are the feasible points.

Table 1: Quantum sensing topics, mathematical toolkits, and selected references.

Topic	Mathematical Toolkit	Selected References
Probe state design under decoherence	SDP, Lyapunov–Schur structure	Demkowicz-Dobrzański et al. [2012]; Liu et al. [2024]; Siddhu and Tayur [2022]
Optimal measurement (POVM) design	SDP, multiparameter Cramér–Rao bound	Hayashi and Ouyang [2023]; Tavakoli et al. [2024]; Siddhu and Tayur [2022]
SDP algorithm design (tailored solvers)	First-order methods, symmetry reduction, warm-starting	Peña et al. [2025]; Ben-Tal et al. [2009]
Adaptive measurement (bandit layer)	Multi-armed bandits, Thompson sampling, Bayesian opt.	Russo et al. [2018]; Simchi-Levi and Wang [2024]; Tang et al. [2025a]
Pulse-sequence design	RL (MDP, policy gradient)	Krenn et al. [2023]; Powell [2022]; Wang et al. [2024]
QFI & QCRB	Quantum information theory, SDP duality	Giovannetti et al. [2011]; Degen et al. [2017]; Liu et al. [2024]
ML surrogates / simulation-based inference	Neural posterior estimation, normalizing flows	Cranmer et al. [2020]; Gan et al. [2025]
LiDAR inverse problem (3-D reconstruction)	TV regularization, robust minimax, compressed sensing	Ben-Tal et al. [2009]; Birge and Louveaux [2011]; Shapiro et al. [2009]
Source attribution (advection-diffusion)	Robust stochastic programming, PDE-constrained opt.	Shapiro et al. [2009]; Ben-Tal et al. [2009]
Scan/sensor placement	Bayesian experimental design, stochastic programming	Chaloner and Verdinelli [1995]; Birge and Louveaux [2011]
Operational dispatch	Stochastic integer programming, UQ propagation	Tang et al. [2025b]; Birge and Louveaux [2011]
MEG source imaging	Ill-posed linear inverse problems, sparse recovery	Boto et al. [2018]; Ben-Tal et al. [2009]
Networked sensor data assimilation	Distributed inference, consensus optimization, MAB	Powell [2022]; Simchi-Levi and Wang [2024]

Observable quantities correspond to Hermitian operators $A = A^\dagger$, with expected value $\langle A \rangle_\rho = \text{Tr}(\rho A)$ and variance $\text{Var}_\rho(A) = \text{Tr}(\rho A^2) - [\text{Tr}(\rho A)]^2$. A measurement yielding outcome a_i occurs with probability $p_i = \text{Tr}(\rho \Pi_i)$, where Π_i is the projector onto the eigenspace of A with eigenvalue a_i . Once probabilities are defined, the data collected by a quantum sensor is a classical random variable and standard statistical methods apply.

The Pauli basis

Any 2×2 Hermitian matrix—and therefore any observable of a qubit—can be written as a linear combination of four fixed matrices called the Pauli matrices. Think of them as the “coordinate axes” of qubit space: just as any vector in 3-D can be decomposed into x -, y -, and z -components, any qubit interaction decomposes into contributions from σ_x , σ_y , σ_z , and the identity I . The Pauli matrices are

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

together with I span all 2×2 Hermitian operators [Nielsen and Chuang, 2010]. Any physical interaction between a two-level sensor and its environment decomposes into these four components.

2.2 Why Semidefinite Programs Arise Naturally

A semidefinite program (SDP) is a convex optimization problem in which the decision variable is a matrix constrained to be positive semidefinite—meaning all its eigenvalues are non-negative. SDPs generalize linear programs (where the variable is a vector) and can be solved in polynomial time with interior-point methods [Ben-Tal et al., 2009], though worst-case polynomial-time theory does not by itself guarantee practical scalability at the matrix dimensions ($2^n \times 2^n$ for n qubits) that arise for even modest numbers of qubits—see the discussion of tailored solvers in Section 8. The connection to quantum sensing is structural: SDPs arise naturally throughout quantum sensing because the basic objects of quantum mechanics are themselves positive semidefinite matrices. Density operators satisfy $\rho \succeq 0$ and $\text{Tr}(\rho) = 1$, so the space of valid probe states is the intersection of the PSD cone with the affine constraint $\text{Tr}(\rho) = 1$ —precisely the basic feasible geometry of an SDP.

A positive operator-valued measure (POVM) is the most general description of a quantum measurement, generalizing the projective measurement $p_i = \text{Tr}(\rho \Pi_i)$ introduced above. Each element E_i of the POVM corresponds to one possible outcome, and the probability of that outcome is $p_i = \text{Tr}(\rho E_i)$. The POVM generalizes the simpler “projective” measurements (like reading out a spin up or down) to include probabilistic and joint measurements—which become important when estimating multiple parameters simultaneously. Physically realizable measurements—POVMs $\{E_i\}$ —are collections of PSD operators satisfying $E_i \succeq 0$ and $\sum_i E_i = I$; optimizing over measurements is again an optimization over a PSD cone. Many important quantum-estimation bounds and channel-optimization problems can therefore be formulated exactly, or bounded tightly, using semidefinite or related conic programs [Demkowicz-Dobrzański et al., 2012, Hayashi and Ouyang, 2023, Liu et al., 2024]. The precise formulation depends on the estimation setting—the noise model, the number of parameters, and what is held fixed as a resource—so no single SDP characterizes every probe-design or measurement-design problem; these formulations are developed in Sections 2.7–2.8 below.

2.3 The Sensing Hamiltonian: Physical Derivation

The Hamiltonian is the mathematical object that governs how a quantum system evolves in time. For a sensor, it plays the role of the “transfer function”: it encodes how the external stimulus (the magnetic field, the gravitational acceleration, the gas concentration) enters the physics of the qubit. Choosing a probe state that interacts strongly with the Hamiltonian is the central design question, and it is the one that Optimization 1 answers.

A quantum sensor’s total Hamiltonian decomposes as

$$\hat{H}(t) = \hat{H}_0 + \hat{H}_V(t) + \hat{H}_{\text{control}}(t), \quad (1)$$

where $\hat{H}_0 = E_{\text{avg}}I - \frac{\Delta E}{2}\sigma_z$ is the free Hamiltonian, $\hat{H}_V(t)$ encodes the stimulus, and $\hat{H}_{\text{control}}(t)$ implements pulse sequences and gates. Because $\hat{H}_V(t)$ is Hermitian, it decomposes as [Degen et al., 2017]:

$$\hat{H}_V(t) = \frac{\gamma}{2} [\text{Re}[V_{\perp}(t)] \sigma_x + \text{Im}[V_{\perp}(t)] \sigma_y - V_{\parallel}(t) \sigma_z], \quad (2)$$

where γ is the transduction parameter, $V_{\parallel}$ shifts the energy splitting, and $V_{\perp}$ couples the two levels. In the limit $\Delta E \gg |V_{\parallel}|, |V_{\perp}|$ [Degen et al., 2017], the abstract formulation $H_s = \phi G$ follows with $G = -\sigma_z/2$, $\phi = \gamma V_{\parallel}$ for a longitudinal stimulus.

Notation convention. We write ρ_0 for the prepared probe state—the experimenter’s choice, the decision variable in Optimization 1. After signal interaction and decoherence, the state becomes $\rho_\phi = \mathcal{E}_\phi(\rho_0)$; the subscript ϕ signals that this state carries information about the unknown parameter. All optimization is over ρ_0 ; ρ_ϕ is a derived quantity determined by the sensing channel $\mathcal{E}_\phi$, which combines the signal Hamiltonian with the noise map $\mathcal{E}$ (Section 2.7).

For a time-independent Hamiltonian, the time-evolution operator is

$$\hat{U}(0, t) \propto \begin{pmatrix} 1 & 0 \\ 0 & e^{it\Delta E/\hbar} \end{pmatrix}, \quad (3)$$

dropping the global phase. When the stimulus is present, $\Delta E \rightarrow \Delta E + \gamma V_\parallel$, encoding $V_\parallel$ in the relative phase between $|0\rangle$ and $|1\rangle$. Measuring this phase is what a quantum sensor does.

2.4 The Four-Step Sensing Protocol

1. **Initialize.** Apply $\hat{U}_{\text{init}}$ to $|0\rangle$ to produce the desired probe superposition.
2. **Interact.** Evolve for time τ under $H_s = \phi G$, accumulating phase proportional to ϕ .
3. **Transform.** Apply $\hat{U}_{\text{readout}}$ (often $\hat{U}_{\text{init}}^{-1}$) to convert accumulated phase into a measurable population difference.
4. **Measure and repeat.** Perform a projective measurement and repeat M times to estimate ϕ .

Running example: NV-center magnetometry. A nitrogen-vacancy (NV) center in diamond is an electron-spin qubit whose sensing Hamiltonian, in the effective two-level and longitudinal-field limit, is $\hat{H}_s = (\gamma_e B/2)\sigma_z$, a direct instantiation of $H_s = \phi G$ with $\phi = \gamma_e B$ and $G = \sigma_z/2$ (the opposite sign from $G = -\sigma_z/2$ in Section 2.3 is a convention: it can be absorbed into ϕ and leaves all Fisher-information quantities unchanged), where γ_e is expressed in angular-frequency units so that the field-induced phase is $\gamma_e B\tau$ (if γ_e is instead defined through an energy coupling, the phase is $\gamma_e B\tau/\hbar$). For the remainder of the paper we report Fisher information with respect to the physical field B itself—the quantity of practical interest—rather than the abstract phase ϕ ; the two are related by $F_Q(\phi) = F_Q(B)/\gamma_e^2$. The four steps above become: (1) prepare $(|0\rangle + |1\rangle)/\sqrt{2}$ with a Hadamard gate; (2) allow the field B to accumulate phase for time τ ; (3) apply a second Hadamard; (4) read out the spin state optically. This protocol is called Ramsey interferometry (Appendix B). The NV center serves as the running example throughout this paper: in the SDP formulations (Sections 2.7–2.8), in the algorithms (Section 4), and in the benchmarking decomposition (Section 7). Table 2 at the end of this section maps the NV example to each layer of Figure 1.

For a single NV center operated in the effective two-level, longitudinal-field regime, the distinctions among separability, generator additivity, and inter-probe locality are absent, so the single-qubit model provides an appropriate effective description. For an NV array, however, these notions are distinct: separability describes the probe state ρ_0 ; additivity describes the generator ($G = \sum_i g_i$); locality describes the coupling or measurement. In an imaging or ensemble geometry, the effective terms g_i can vary with probe–sample distance, dipolar coupling kernels, and field gradients; probe–probe interactions and correlated environmental noise can further invalidate an independent-probe model. Additivity of G alone does not deliver the Heisenberg-limit scaling of Section 2.5—even with local interactions among array elements, a separable initial state cannot beat the standard quantum limit [Shi et al., 2024]. Reaching the Heisenberg limit requires genuine entanglement in ρ_0 , the GHZ (Greenberger–Horne–Zeilinger) construction of Appendix B, not merely a sum-of-local-terms generator. For an NV ensemble, the independent-probe SQL benchmark further assumes that the

individual couplings are effectively local and that preparation, control, and readout do not introduce collective correlations; if only collective control or aggregate fluorescence readout is available, the achievable Fisher information must be evaluated for the corresponding global measurement model.

2.5 Key Quantum Limits

Two precision limits bound every protocol. With N independent probes and M shots, the *standard quantum limit* (SQL) is $\widehat{\Delta\phi}_{\text{SQL}} = 1/\sqrt{NM}$. Entangling all N probes in a GHZ state yields a quadratic reduction in the required number of probes for a fixed target precision (equivalently, a $\sqrt{N}$ reduction in standard deviation at fixed N), reaching the *Heisenberg limit* (HL): $\widehat{\Delta\phi}_{\text{HL}} = 1/(N\sqrt{M})$ [Giovannetti et al., 2011]. For the NV running example with a single qubit ($N = 1$), precision is governed directly by the single-probe QCRB; SQL and HL become distinct particle-number scalings only when $N > 1$. In an ideal noiseless model with a suitably additive generator and collective state preparation/readout, entanglement across an NV array can attain HL scaling; decoherence and correlated noise can reduce or eliminate that asymptotic advantage. Full derivations, including the inseparability proof for the GHZ state and the Ramsey Fisher-information calculation, are given in Appendix B.

2.6 Quantum Fisher Information and the Cramér–Rao Bound

The *quantum Fisher information* (QFI) $F_Q(\rho_\phi)$ is the supremum of the classical Fisher information F_C over all POVMs, where for a POVM $\{E_i\}$ giving outcome probabilities $p_i(\phi) = \text{Tr}(\rho_\phi E_i)$,

$$F_C = \sum_i \frac{[\partial_\phi p_i(\phi)]^2}{p_i(\phi)}$$

is the same quantity familiar from classical statistics: it measures how sharply the outcome distribution changes with ϕ and, through the classical Cramér–Rao bound, sets a lower bound on the variance of any locally unbiased estimator constructed from the data. The *quantum Cramér–Rao bound* (QCRB) states that for any **locally unbiased estimator of a scalar parameter ϕ** based on M measurements:

$$\text{Var}(\hat{\phi}) \geq \frac{1}{MF_Q(\rho_\phi)}. \quad (4)$$

For a *pure* probe state undergoing noiseless encoding via $H_s = \phi G$ over interrogation time τ (the Four-Step Protocol above), the QFI equals four times the variance of the accumulated-phase generator τG [Degen et al., 2017]: $F_Q = 4\tau^2 \text{Var}_{\psi_0}(G)$. (If instead ϕ is defined to directly parameterize the unitary $U_\phi = e^{-i\phi G}$ with no separate interrogation time, drop the τ^2 factor; the two conventions differ only in whether τ is absorbed into G or left explicit, and this paper uses the explicit- τ convention throughout.)

For a mixed state $\rho = \sum_m p_m |m\rangle\langle m|$ undergoing encoding via $H_s = \phi G$ over interrogation time τ [Degen et al., 2017], the general definition is

$$F_Q = 2\tau^2 \sum_{\substack{m,n: \\ p_m + p_n > 0}} \frac{(p_m - p_n)^2}{p_m + p_n} |\langle m|G|n\rangle|^2. \quad (5)$$

Section 2.7 develops an equivalent route to this quantity via the symmetric logarithmic derivative (SLD) operator. Here p_m and $|m\rangle$ are the eigenvalues and eigenvectors of ρ . Each term weights the coupling between eigenstates m and n by how distinguishable their populations are: the weight

$(p_m - p_n)^2/(p_m + p_n)$ vanishes when $p_m = p_n$ (indistinguishable populations carry no information) and grows as one population dominates the other. As a check, when ρ is pure (one $p_m = 1$, all others zero), Eq. (5) reduces to the familiar $F_Q = 4\tau^2 \text{Var}_{\psi_0}(G)$ above. In general $F_Q[\rho, G] \leq 4\tau^2 \text{Var}_\rho(G)$, with the inequality generally strict: the total variance can include classical statistical or thermal fluctuations that carry no useful quantum coherence for parameter estimation [Hauke et al., 2016]. Real sensors are typically mixed due to finite temperature, imperfect preparation, and decoherence—the case treated immediately below. QFI has been measured experimentally in solid-state qubits: Yu et al. [2022] extracted it from coherent dynamical response in an NV-center spin and demonstrated near saturation of the QCRB in phase estimation.

For the NV running example under pure dephasing at rate Γ , $F_Q = (\gamma_e \tau)^2 e^{-2\Gamma\tau}$, which falls below the noiseless value $F_Q^0 = (\gamma_e \tau)^2$ —the decoherence loss discussed in Section 7. Once decoherence is present, the pure-state variance formula no longer applies and the optimization formulation of the next subsection becomes essential.

2.7 Optimization 1: Designing the Optimal Probe State

Question answered: given the specific noise model $\mathcal{E}$ of a physical sensor and the parameter ϕ encoded by generator G , what is the best quantum state to prepare?

Throughout, $\mathcal{E}$ denotes the parameter-independent noise (decoherence) map of the hardware, and $\mathcal{E}_\phi$ denotes the parameter-dependent *sensing channel*: signal encoding by $H_s = \phi G$ over interrogation time τ together with the noise acting during that interval. Both are completely positive, trace-preserving (CPTP) linear maps, and the probe state after interrogation is $\rho_\phi = \mathcal{E}_\phi(\rho_0)$, which is mixed whenever $\mathcal{E}$ is non-trivial. When the noise commutes with the encoding—as for the pure dephasing in the NV example below—the sensing channel factorizes as

$$\mathcal{E}_\phi(\rho_0) = \mathcal{E}(U_\phi \rho_0 U_\phi^\dagger) = U_\phi \mathcal{E}(\rho_0) U_\phi^\dagger, \quad U_\phi = e^{-i\phi\tau G},$$

so Eq. (5) applies directly with $\rho = \mathcal{E}(\rho_0)$. For general noise acting concurrently with the signal, $\mathcal{E}_\phi$ is the map generated by the corresponding master equation over $[0, \tau]$ and need not factorize. Only $\mathcal{E}_\phi$ carries information about ϕ ; $\mathcal{E}$ determines how much of it survives.

The CPTP condition is the quantum analogue of a stochastic matrix: it maps valid quantum states to valid quantum states, preserving non-negativity and normalization. The symmetric logarithmic derivative (SLD) operator L is the quantum analogue of the classical score function: it satisfies $\partial_\phi \rho_\phi = (\rho_\phi L + L \rho_\phi)/2$, and the QFI is $F_Q = \text{Tr}(\rho_\phi L^2)$. This is the same quantity as Eq. (5) of Section 2.6, now expressed through the SLD rather than ρ 's eigendecomposition directly.

Given the sensing channel $\mathcal{E}_\phi$ defined above, and with F_Q evaluated at the true value (or operating point) of ϕ , the probe-design problem is

$$F_Q^* = \max_{\rho_0 \succeq 0, \text{Tr}(\rho_0)=1} F_Q[\mathcal{E}_\phi(\rho_0)]. \quad (6)$$

This compact expression states the physical design problem: choose the input state that makes the channel output maximally informative about ϕ . Depending on the channel model and the allowed sensing strategy, this optimization may admit exact SDP formulations or SDP-computable bounds [Demkowicz-Dobrzański et al., 2012, Liu et al., 2024, Hayashi and Ouyang, 2023]; the formulation is more elaborate than a single generic block-matrix program and is specific to the noise model at hand. What matters for the optimization–inference stack developed here is not a particular SDP representation but the certified value F_Q^* and, when recoverable, an optimal probe state ρ_0^* .

NV running example. For the NV center under pure dephasing at rate Γ , the noise map $\mathcal{E}$ shrinks the off-diagonal elements of ρ_0 by $e^{-\Gamma\tau}$. **Because this dephasing commutes with $G \propto \sigma_z$, $\mathcal{E}_\phi$ factorizes as above.** Direct evaluation of Eq. (5) shows that the Hadamard-prepared state $(|0\rangle + |1\rangle)/\sqrt{2}$ maximizes the QFI over this one-qubit dephasing channel: state-preparation loss is zero, and the achievable $F_Q^* = (\gamma_e\tau)^2 e^{-2\Gamma\tau}$ equals the physical decoherence ceiling.

Take-home for non-quantum readers. A practitioner who solves Eq. (6)—whether analytically, as in the NV example, or numerically via an SDP/conic-program formulation appropriate to the noise model—receives two outputs: the optimal probe state ρ_0^* to prepare (an experimental instruction) and the value F_Q^* (the precision ceiling no downstream algorithm can breach for this hardware and noise model). The formulation is the means; the output is what matters operationally.

2.8 Optimization 2: Designing the Measurement

Question answered: given the optimal probe state, what measurements should the experimenter perform?

We write $\phi \in \mathbb{R}$ (regular weight) for a single scalar parameter, as in the preceding subsections, and $\boldsymbol{\phi} = (\phi_1, \dots, \phi_d)^\top \in \mathbb{R}^d$ (boldface) for a vector of d parameters estimated simultaneously (e.g., $d = 3$ for the three components of a magnetic field); its components ϕ_k are scalars and are therefore set in regular weight. Boldface is used for every vector-valued quantity in the remainder of this subsection. Given a probe state ρ_ϕ , measurement design chooses a POVM $\{E_i\}$ to maximize the information extracted about the unknown parameter vector $\boldsymbol{\phi}$. Eq. (7) is stated for general d ; for $d = 1$, $\boldsymbol{\phi}$ reduces to the scalar ϕ , Σ to $\text{Var}(\hat{\phi})$, and W to a positive scalar. Concretely, the measurement-design problem is

$$\min_{\{E_i\}, \hat{\phi}} \text{Tr}[W \Sigma_\phi(\{E_i\}, \hat{\phi})] \quad \text{s.t.} \quad E_i \succeq 0, \sum_i E_i = I, \quad \hat{\phi} \text{ locally unbiased at } \boldsymbol{\phi}, \quad (7)$$

where $p(i | \boldsymbol{\phi}) = \text{Tr}[\rho_\phi E_i]$ is the outcome likelihood induced by the POVM, $\hat{\phi} : i \mapsto \hat{\phi}(i) \in \mathbb{R}^d$ is an estimator assigning a parameter estimate to each outcome, and

$$\Sigma_\phi(\{E_i\}, \hat{\phi}) = \sum_i p(i | \boldsymbol{\phi}) (\hat{\phi}(i) - \boldsymbol{\phi})(\hat{\phi}(i) - \boldsymbol{\phi})^\top$$

is the single-shot covariance they jointly induce (with M independent repetitions it scales as Σ/M). Local unbiasedness at $\boldsymbol{\phi}$ requires $\sum_i p(i | \boldsymbol{\phi}) \hat{\phi}_k(i) = \phi_k$ and $\sum_i \partial_{\phi_j} p(i | \boldsymbol{\phi}) \hat{\phi}_k(i) = \delta_{jk}$ for $j, k = 1, \dots, d$. This is the condition under which the classical and quantum Cramér–Rao bounds apply, and it excludes trivial estimators that achieve zero variance by ignoring the data. $W \in \mathbb{S}_{++}^d$ is a fixed, application-specific weight matrix. Σ is *not* an independent decision variable: it is determined by the choice of $\{E_i\}$ and $\hat{\phi}$, which are the only quantities optimized. For a fixed POVM with invertible classical Fisher information matrix $F_C(\boldsymbol{\phi})$, the inner minimum over locally unbiased estimators equals $\text{Tr}[W F_C(\boldsymbol{\phi})^{-1}]$, so Eq. (7) is equivalently $\min_{\{E_i\}} \text{Tr}[W F_C^{-1}]$. This problem is generally nonlinear in $\{E_i\}$; it is a conceptual formulation of the design goal, not itself a displayed SDP.

For single-parameter estimation, measurement in the eigenbasis of the SLD attains the QFI under the usual regularity conditions, giving an explicit optimal $\{E_i\}$ for Eq. (7) in that case. For multiparameter estimation—simultaneously estimating all three spatial components of a magnetic field, for example—the optimal measurements for different parameters may be incompatible, since no single POVM achieves all bounds simultaneously in general [Hayashi and Ouyang, 2023, Liu

et al., 2024]. In finite-dimensional multiparameter models, the Holevo Cramér–Rao bound provides an SDP-computable *lower bound* on the attainable weighted covariance in Eq. (7) [Hayashi and Ouyang, 2023, Liu et al., 2024]; recovering a physically implementable POVM that attains it may require a separate, model-specific construction.

Decision variables—physical meaning.

- $\{E_i\}$: the POVM—which outcomes are possible and with what probabilities.
- $\hat{\phi}$: the estimator—how recorded outcomes are mapped to parameter estimates.

Induced quantity and weights. Σ is the estimator covariance induced by $(\{E_i\}, \hat{\phi})$ —the error budget—scalarized by an application-specific matrix $W \succ 0$. For navigation, position errors may be weighted more heavily than orientation; for magnetoencephalography (MEG), the direction of maximum neuronal activity dominates.

Optimizations 1 and 2 together form a two-stage design pipeline. The Fisher-information gap between a practical protocol and the relevant bound decomposes into state-preparation loss (closeable by Optimization 1) and measurement loss (closeable by Optimization 2).

Take-home for non-quantum readers. A practitioner who solves the measurement-design problem receives an optimal POVM and its induced covariance when an attaining construction is available. In multiparameter settings, an HCRB/SDP evaluation instead provides a certified lower bound on the weighted covariance; constructing an attaining measurement may require a separate, model-specific step. Together, Optimization 1 and Optimization 2 serve as design tools that indicate whether a given experimental protocol is losing precision due to a bad probe state, a bad measurement choice, or unavoidable decoherence—and the answer directly determines where engineering effort should be focused.

NV running example (continued). For the NV Ramsey protocol, an appropriate final readout rotation followed by σ_z projective measurement attains $F_C = F_Q$ at the optimal operating point; hence measurement loss is zero in this idealized single-parameter example. The full precision gap is therefore pure decoherence loss, quantified exactly by the inequality chain in Section 7.

3 Hardware Platforms

Hardware diversity in quantum sensing is not incidental. Each platform occupies a distinct position in the space of noise geometry, control cost, and readout structure—and each position implies a different set of algorithmic bottlenecks. Nitrogen-vacancy (NV) center arrays in diamond experience predominantly dephasing noise with long coherence times, room-temperature operation, and optical readout that is efficient but photon-starved; their bottleneck is probe-state design under correlated noise, a natural SDP problem. Superconducting quantum interference devices (SQUIDs) experience low-frequency flux noise and require cryogenic infrastructure, but their readout is fast and high-fidelity; their bottleneck is source reconstruction from a limited sensor array, a classical inverse problem. Atom interferometers are exquisitely sensitive to vibration; their bottleneck is robust estimation under non-stationary disturbance, a stochastic filtering problem. Single-photon LiDAR is Poisson-limited and ill-conditioned in the spatial reconstruction; its bottleneck is compressed sensing and source attribution under meteorological uncertainty. Table 3 summarizes the platforms discussed in this section.

Table 2: NV-center magnetometry running example mapped to the optimization–inference stack of Figure 1. Each row is one pipeline stage; the bounds column records the output of Optimization 1 or Optimization 2 for this noise model. All losses in the classical-layer column are Fisher-information losses, defined in Section 7 (Eq. (26)). The decoherence loss in the final row is irreducible for this noise model and is the only non-zero gap component, as established in Section 7.

Stage	Quantum layer	Bounds layer	Classical layer
State design	Prepare $(0\rangle+ 1\rangle)/\sqrt{2}$ via Hadamard gate	Optimization 1 (analytic) identifies the Hadamard state as optimal	State-preparation loss = 0
Signal encoding	Field B imprints phase $\gamma_e B \tau$ (Ramsey protocol)	$F_Q^* = (\gamma_e \tau)^2 e^{-2\Gamma \tau}$ (Optimization 1 output)	QCRB floor: $\text{Var}(\hat{B}) \geq 1/(MF_Q^*)$
Measurement	Readout rotation + σ_z projective readout via optical spin detection	SLD-eigenbasis argument (Optimization 2) confirms this attains $F_C = F_Q$	Measurement loss = 0
Pulse design	Dynamical decoupling suppresses dephasing rate Γ	RL optimality gap δ_{RL} benchmarked against F_Q^*	RL agent selects next pulse type and timing
Decision	Raw spin-flip probability from photon counts	Decoherence loss $1 - e^{-2\Gamma \tau}$ (irreducible)	Posterior over B ; report field estimate $\hat{B}$

3.1 Magnetic Sensing

NV centers in diamond [Lab/prototype] are point defects in which a nitrogen atom substitutes for a carbon atom adjacent to a lattice vacancy. The electron spin is a qubit initialized and read out optically at room temperature, with sensing Hamiltonian, in the effective two-level and longitudinal-field limit, $\hat{H}_s = (\gamma_e B/2)\sigma_z$ —a direct instantiation of $H_s = \phi G$ with $\phi = \gamma_e B$ and $G = \sigma_z/2$. Coherence times reach milliseconds in isotopically purified diamond [Barry et al., 2020].

SQUIDs [Deployed] exploit the Josephson effect and magnetic flux quantization to achieve femtotesla sensitivity. Clinical SQUID-based MEG systems have been in routine use for epilepsy presurgical evaluation since the 1990s; cryogenic cooling to ~ 4 K constrains installation to fixed clinical facilities.

OPMs [Field trial] use laser light to align atomic spins in an alkali vapor cell, achieving femtotesla sensitivity at room temperature. Boto et al. [2018] demonstrated a wearable OPM MEG system tolerating natural head movement; several systems are now in clinical evaluation.

Local imaging versus momentum-resolved spectroscopy

NV-based imaging resolves local fields, defects, interfaces, domains, and spatial inhomogeneity, but the measured signal is filtered by the probe geometry and its spatial coupling kernel—the locality issue raised in Section 2.3. Inelastic neutron and x-ray scattering take a fundamentally different approach: a beam of neutrons or x-rays is fired at the sample, and the momentum and energy the beam loses (or gains) on scattering is recorded. Because the beam interacts coherently with the entire illuminated volume at once, rather than with one probe atom at a time, this

Table 3: Quantum sensing hardware platforms. fT = femtotesla; pT = picotesla; ppm·m = parts per million times meters (path-integrated concentration); MEG = magnetoencephalography; TRL = Technology Readiness Level.

Platform	Modality	Sensitivity	Cryo?	TRL	Maturity
NV center (diamond)	Magnetic, thermal	fT-pT Hz ^{-1/2}	No	5–7	[Lab/prototype]
SQUID	Magnetic	< 1 fT Hz ^{-1/2}	Yes	8–9	[Deployed]
OPM	Magnetic	~fT Hz ^{-1/2}	No	7–8	[Field trial]
Atom interferometer	Gravity, rotation	10 ⁻⁸ m s ⁻²	No	6–8	[Field trial]
Atomic clock (microwave)	Frequency	~ 10 ⁻¹⁶	No	8–9	[Deployed]
Optical lattice clock	Frequency	< 10 ⁻¹⁸	No	7–8	[Lab/prototype]
Single-photon LiDAR	Range, gas conc.	ppm·m at 10–250 m	No	6–8	[Field trial]

Table 4: Local imaging (NV) versus momentum-resolved spectroscopy (neutron/x-ray) as complementary quantum sensing modalities.

Attribute	NV imaging (local)	Neutron/x-ray (momentum-resolved)
Generator support	Local, spatially disjoint per site	Extended — $G_q = \text{FT}[O_j]$
Spatial resolution	High (nanometer-scale, direct)	Low (averaged over sample)
Momentum resolution	None	High — direct access to q, ω
Sample volume probed	Local (near-field, μm –mm)	Macroscopic (bulk)
Disorder / heterogeneity sensitivity	High — sees local defects, domains directly	Low — averages over disorder
Accessible correlation structure	Single-site / local field values	Collective modes, non-local correlations

reveals *collective* excitations—correlated, wave-like motion shared across many particles (a spin wave rippling through a magnet, or a lattice vibration)—rather than the state of any single site. Formally, this probes generators of the form $G_q = \text{FT}[O_j]$ (the Fourier transform of a local operator) that remain additive in local operators but have extended spatial support and select collective modes at a particular momentum q and energy ω (the momentum and energy transferred between probe and sample). These reciprocal-space measurements provide access to non-local correlations and momentum-dependent QFI that are difficult to reconstruct from local imaging alone; on the other hand, neutron and x-ray measurements typically average over a macroscopic sample and can hide disorder, domains, and spatial heterogeneity that are directly visible to local probes. The two approaches are therefore complementary rather than interchangeable: local imaging identifies *where* the response occurs, while spectroscopy identifies its momentum, energy, and collective correlation structure [Mathew et al., 2020, Scheie et al., 2021, Shen et al., 2026]; [Table 4 summarizes the contrast](#). In the optimization–inference language of this paper, the two techniques instantiate the same stack with a different generator and a different measurement space: probe design for a scattering experiment means choosing which incident beam momenta and energies to scan—the analogue of choosing ρ_0 in Section 2.7—and measurement design means choosing detector placement and energy windows, the analogue of choosing a POVM in Section 2.8. Inverting the resulting noisy photon or neutron counts into a momentum-energy spectrum is again an inverse problem of the kind Section 6 develops for methane LiDAR, just with q and ω in place of range and concentration.

3.2 Inertial and Gravitational Sensing

Atom interferometers [Field trial] accumulate differential phase under gravity or rotation. According to the manufacturer, the Muquans/Exail Absolute Quantum Gravimeter (AQG) has been deployed in field campaigns for hydrology and geophysical monitoring [Bongs et al., 2019, Muquans/Exail, 2021]. Quantum inertial navigation systems (Q-INS) are under government-funded development in the US and UK [Bongs et al., 2019] but remain in early field-trial or laboratory stages.

3.3 Timing and Frequency

Microwave atomic clocks [Deployed] underpin the GPS constellation, financial timestamping, and power grid synchronization [NIST, 2025]—the most mature quantum sensing platform.

Optical lattice clocks [Lab/prototype] have achieved fractional frequency uncertainties below 10^{-18} in laboratory demonstrations and remain laboratory instruments.

3.4 Photonic Sensing: Single-Photon LiDAR

Single-photon LiDAR [Field trial] combines single-photon avalanche **diodes** (SPADs) with time-correlated single-photon counting (TCSPC) and differential absorption LiDAR (DIAL) to achieve standoff gas concentration imaging. QLM Technology reports having received US EPA Alternative Test Method qualification in 2025 [QLM Technology, 2025]; Bridger Photonics deploys an airborne system for pipeline inspection, as reported by MIT News [MIT News Office, 2025].

4 Sensing Algorithms and Signal Processing

The two sensing-design optimizations above, together with their SDP/conic bounds where available, set the precision ceiling. The algorithms in this section are what actually run—in the sensor, in the cloud, and in the decision system. They range from classical Bayesian updating to bandit protocols to reinforcement learning, and each is given a formal mathematical specification commensurate with the optimization formulations of Section 2.

4.1 Bayesian Estimation and Adaptive Protocols (Bandits)

Bayesian inference is the natural framework for quantum sensing because the protocol defines an explicit likelihood. Given unknown parameter ϕ and measurement setting ξ (pulse duration, laser wavelength, measurement axis), the outcome x is drawn from a known distribution $p(x|\phi, \xi)$. The posterior evolves by standard Bayesian updating:

$$p(\phi \mid x_{1:t}) \propto p(x_t \mid \phi, \xi_t) p(\phi \mid x_{1:t-1}). \quad (8)$$

Adaptive sensing as a bandit problem

In a multi-armed bandit problem, an agent repeatedly chooses among several “arms” (actions), receives a reward, and must balance exploring new arms against exploiting the arm that seems best so far [Russo et al., 2018]. In adaptive quantum sensing, each “arm” is a measurement setting ξ (e.g., the duration of a Ramsey sequence or the orientation of a probe field), and the “reward” is the information gained about ϕ . The coherence time T_2 bounds each individual interrogation, and the

mission time T_{mission} bounds the total run: once the qubit decoheres within a shot, or the mission budget is exhausted, no further useful measurements can be made.

After each measurement, the next setting ξ_{t+1} is chosen to maximize expected information gain (IG) over the current posterior. This is exactly the explore–exploit trade-off of a multi-armed bandit (MAB) [Russo et al., 2018, Simchi-Levi and Wang, 2024]. Formally, the adaptive sensing problem is:

$$\max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=1}^M \text{ig}(\xi_t, x_t) \right] \quad \text{subject to} \quad \tau(\xi_t) \leq T_2 \quad \forall t, \quad \sum_{t=1}^M [\tau(\xi_t) + t_{\text{overhead}}] \leq T_{\text{mission}}, \quad (9)$$

where T_2 is the coherence time bounding a single interrogation and T_{mission} is the aggregate time budget for the run, including per-shot overhead t_{overhead} ; the per-step realized information gain is

$$\text{ig}(\xi_t, x_t) = H[p(\phi \mid x_{1:t-1})] - H[p(\phi \mid x_{1:t})], \quad (10)$$

and the expected information gain used to select ξ_t *before* observing x_t is

$$\text{IG}(\xi_t) = H[p(\phi \mid x_{1:t-1})] - \mathbb{E}_{x_t \mid \xi_t, x_{1:t-1}} [H[p(\phi \mid x_{1:t})]]. \quad (11)$$

$H[\cdot]$ denotes Shannon entropy and $\tau(\xi_t)$ is the duration of measurement ξ_t .

The bounds layer enters as an upper bound on per-step gain, under a local Gaussian/Laplace approximation to the posterior update: Fisher information does not in general upper-bound mutual information for arbitrary likelihoods and priors, so the following is an approximation valid near a well-concentrated posterior, not a general inequality. For a locally Gaussian posterior approximation with variance σ_t^2 , the expected information gain satisfies [Giovannetti et al., 2011, Russo et al., 2018]:

$$\text{IG}(\xi_t) \lesssim \frac{1}{2} \log(1 + \sigma_t^2 F_C(\xi_t)) \leq \frac{1}{2} \log(1 + \sigma_t^2 F_Q^*(\xi_t)), \quad (12)$$

where $F_Q^*(\xi_t)$ is the optimal QFI for measurement setting ξ_t (Optimization 1); **the bound is** increasing in the prior variance σ_t^2 as expected: more uncertainty leaves more information to be gained.

Thompson sampling [Russo et al., 2018] applies with a sampled-utility rule: draw a posterior sample and choose the setting that maximizes a utility function evaluated at that sample, e.g. the classical Fisher information it would achieve:

$$\tilde{\phi}_t \sim p(\phi \mid x_{1:t-1}), \quad \xi_t \in \arg \max_{\xi \in \Xi} u(\xi; \tilde{\phi}_t), \quad u(\xi; \tilde{\phi}_t) = F_C(\tilde{\phi}_t, \xi), \quad (13)$$

or, more generally, an application-specific value-of-information function in place of F_C .

This bandit formulation is structurally analogous to problems the INFORMS data science community has already solved. Multi-armed bandits with endogenous learning curves have been used for split liver transplantation allocation [Tang et al., 2025a]; in quantum sensing the analogous endogeneity is that each measurement updates the posterior over ϕ , changing the optimal next measurement setting via (13). The Pareto-optimal trade-off between regret and statistical power of Simchi-Levi and Wang [2024] maps directly onto the precision-versus-coherence-budget trade-off in (9). Active sequential hypothesis testing under assay budget constraints [Gan et al., 2025] shares the same sequential structure as adaptive quantum sensing.

ML surrogates for Bayesian updating

At high shot rates, sequential Bayesian updating via particle filters is computationally costly. Amortized inference addresses this: a neural network q_θ trained offline on simulated $(\phi, x_{1:t})$ pairs outputs the posterior parameters given observed outcomes. The training objective is:

$$\hat{\theta} = \arg \min_{\theta} \mathbb{E}_{(\phi, x_{1:t}) \sim p} [D_{\text{KL}}(p(\phi | x_{1:t}) \| q_\theta(\phi | x_{1:t}))]. \quad (14)$$

The QFI bound is most naturally used to benchmark a specified point estimator on simulated or controlled data, not to certify a Bayesian posterior directly. In a local regime where the relevant Cramér–Rao conditions apply, define the efficiency ratio

$$\eta_{\text{QFI}} = \frac{1/(MF_Q^*)}{\text{MSE}(\hat{\phi})} \leq 1. \quad (15)$$

Posterior surrogates should instead be evaluated using calibration, coverage, proper scoring rules, and posterior predictive checks [Cranmer et al., 2020]; posterior variance alone is not a QCRB efficiency certificate, since an informative prior can shrink it below $1/(MF_Q^*)$ while a broad prior, finite-sample nonlocality, model mismatch, or a non-Gaussian posterior can inflate it above that value without indicating a suboptimal surrogate.

4.2 Reinforcement Learning for Pulse-Sequence Design

Quantum coherence is fragile: environmental noise (phonons, nearby nuclear spins, electromagnetic fluctuations) causes the qubit to lose its phase relationship over time, a process called decoherence. Dynamical decoupling suppresses this by applying carefully timed microwave pulses that periodically flip the qubit spin, effectively averaging out low-frequency noise. The design question—which pulse, when, for how long—is a sequential decision problem under uncertainty, which maps naturally onto a Markov decision process (MDP).

This is the most direct RL application in quantum sensing: RL agents have been reported to discover pulse sequences that outperform analytically designed protocols [Krenn et al., 2023]. The design of optimal pulse sequences under constraints on total sensing time and pulse power is formalized as an MDP [Powell, 2022]:

- **State:** $s_t = (\rho_t, \Delta t_{\text{rem}}, \hat{\mathcal{E}}_t)$ —current density matrix, remaining time budget, and **estimated noise map**.
- **Actions:** a_t —choice of next pulse (type, timing, duration).
- **Reward:** terminal reward tied to the SDP optimality gap: $r(s_t, a_t) = F_Q(\rho_T)/F_Q^*$ at $t = T$, zero otherwise.
- **Horizon:** the coherence time T_2 / total experiment time.

NV running example (RL layer). For the NV center under simple pure dephasing, the Ramsey protocol is already optimal and $\delta_{\text{RL}} = 0$. Under more complex or non-stationary noise—realistic in a laboratory environment—RL discovers pulse sequences (such as modified XY-8 variants—a family of periodic pulse sequences that refocus dephasing noise) that reduce δ_{RL} below what the analytically designed CPMG sequence (Carr–Purcell–Meiboom–Gill, a standard dynamical-decoupling pulse sequence) achieves [Krenn et al., 2023]. Policy gradient optimization via REINFORCE—a standard

Monte Carlo policy-gradient algorithm—yields:

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\pi}[\nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \cdot R_t], \quad R_t = \sum_{k=t}^T \gamma_d^{k-t} r_k. \quad (16)$$

where R_t is the discounted return and $\gamma_d \in (0, 1]$ the discount factor (distinct from the generator G and the transduction parameter γ of Section 2.3). The *RL optimality gap* is:

$$\delta_{\text{RL}} = 1 - \frac{F_Q(\rho_T)}{F_Q^*}. \quad (17)$$

Off-policy evaluation methods [Wang et al., 2024] provide rigorous uncertainty quantification for this gap when the deployment noise spectrum differs from training. Krenn et al. [2023] demonstrated that RL agents discover pulse sequences that outperform XY-8 and CPMG under realistic noise spectra, particularly for non-stationary or non-Gaussian noise that analytical designs cannot handle.

4.3 Point Cloud Processing and Inverse Problems for Single-Photon LiDAR

A single-photon LiDAR gas scan produces, for each pixel, a time-resolved photon arrival histogram. Converting these into actionable gas concentration estimates involves three algorithmic stages.

Stage 1: Range and absorption extraction. Fit a Poisson model to the TCSPC histogram to extract time of flight and differential absorption depth. In low-photon-count scenarios this is a regularized maximum-likelihood problem:

$$\hat{\mu} = \arg \min_{\mu \geq 0} \sum_{i=1}^B [\mu(r_i) - n_i \log \mu(r_i)] + \lambda_1 \|\mu\|_{\text{TV}}, \quad (18)$$

where, for a given pixel and laser wavelength, the TCSPC histogram has B time bins with centers t_i ; $r_i = c_{\text{light}} t_i / 2$ is the corresponding range; n_i is the photon count in bin i ; and $\mu : r \mapsto \mu(r) \geq 0$ is the intensity function—the expected number of detected photons per bin at range r (target return plus background)—so that $n_i \sim \text{Poisson}(\mu(r_i))$ independently across bins. The optimization is over the vector $(\mu(r_1), \dots, \mu(r_B))$, and $\|\mu\|_{\text{TV}} = \sum_{i=1}^{B-1} |\mu(r_{i+1}) - \mu(r_i)|$ is its (discrete) total variation seminorm. The regularization parameter $\lambda_1 > 0$ trades data fidelity against roughness and is chosen in practice by, e.g., the discrepancy principle, the L-curve, or cross-validation. The bracketed term is the negative Poisson log-likelihood (up to an additive constant), so the reconstruction is a minimization: the TV penalty discourages roughness rather than favoring it. Time of flight is read from the location of the return peak in $\hat{\mu}$, and the differential absorption depth from the on- and off-line return intensities via Eq. (21).

Stage 2: Three-dimensional concentration reconstruction. Line-integrated measurements $C_i = \int_{\ell_i} c(\mathbf{x}) dl$ define the linear inverse problem $C = Kc$, where K is a tall, sparse, ill-conditioned ray-integration matrix. The TV-regularized reconstruction is:

$$\hat{c} = \arg \min_{c \geq 0} \|Kc - C\|_2^2 + \lambda_2 \|\nabla c\|_1. \quad (19)$$

Here $c \in \mathbb{R}_{\geq 0}^V$ collects methane concentrations on V voxels, $C \in \mathbb{R}^P$ the P path-integrated measurements, $K \in \mathbb{R}^{P \times V}$ has entries K_{iv} equal to the length of ray ℓ_i within voxel v , ∇ is the discrete gradient on the voxel grid, and $\lambda_2 > 0$ is a regularization parameter (chosen as for λ_1).

Stage 3: Source attribution and flux estimation. Fitting an advection-diffusion model to the reconstructed field localizes emission sources. Wind speed and direction are uncertain inputs, leading to the robust minimax formulation [Ben-Tal et al., 2009]:

$$\hat{S} = \arg \min_{S \geq 0} \max_{u \in \mathcal{U}} \|\mathcal{K}(u)S - C\|_2^2 + \lambda_3 \|S\|_1, \quad (20)$$

where $S \in \mathbb{R}_{\geq 0}^J$ is the vector of emission rates (e.g., kg h^{-1}) at J candidate source cells—a discretization of the source term $S(\mathbf{x})$ in Eq. (23)—and u is the wind (speed and direction, or a discretized wind field $\mathbf{u}$). The wind-dependent forward operator $\mathcal{K}(u) = K A(u) \in \mathbb{R}^{P \times J}$ maps source rates to predicted path-integrated measurements: $A(u) \in \mathbb{R}^{V \times J}$ is the discretized (e.g., steady-state) solution operator of the advection-diffusion equation (23), whose j th column is the concentration field produced by a unit-rate source at cell j under wind u , and K is the ray-integration matrix of Eq. (19). The uncertainty set is $\mathcal{U} = \{u : (u - \bar{u})^\top P_u^{-1}(u - \bar{u}) \leq 1\}$, an ellipsoid around the measured wind $\bar{u}$ with shape matrix $P_u \succ 0$ reflecting anemometer uncertainty, and $\lambda_3 > 0$ weights the ℓ_1 penalty that promotes a sparse set of active sources (a few leaks).

5 Applications

5.1 Medical Imaging: Magnetoencephalography

SQUID MEG [Deployed] has been used clinically for epilepsy presurgical evaluation for several decades. MEG source imaging—reconstructing the neuronal current distribution from external field measurements—is an ill-posed linear inverse problem [Ben-Tal et al., 2009]. Bayesian inference with structured priors, sparse recovery, and neural surrogates trained on biophysical simulations are active research directions.

Wearable OPM MEG [Field trial] places sensors directly on the scalp, offering higher signal amplitude and tolerance of head movement. Several such systems are reportedly progressing toward clinical evaluation.

5.2 Geophysical Surveying

Atom-interferometer gravimeters [Field trial] detect subsurface density anomalies at economically relevant scales: aquifer volume changes, volcanic intrusion, subsurface voids. Optimal placement of a finite number of gravimeter platforms to maximize information about a subsurface target is a stochastic programming problem [Chaloner and Verdinelli, 1995, Birge and Louveaux, 2011].

5.3 Environmental Emissions Monitoring

Methane LiDAR [Field trial]: methane has a global warming potential approximately 80 times that of CO_2 over a 20-year horizon (Section 3 notes QLM’s reported EPA qualification). Section 6 provides a complete end-to-end treatment of this application, including a discussion of the classical-versus-quantum measurement trade-offs for single-photon detection at eye-safe pulse energies.

5.4 Emerging Directions

Quantum inertial navigation [Lab/prototype] based on atom interferometers is being investigated as a route to lower long-term drift than conventional inertial measurement units. Government-funded field trials are reportedly underway in the US and UK. Atom interferometers also test the weak

equivalence principle at the 10^{-12} level and can detect gravitational waves at sub-Hz frequencies complementary to LIGO. NV-center magnetometers have imaged nuclear spin networks of 50 atoms surrounding a single defect [Barry et al., 2020]. These directions are at laboratory scale and are included to illustrate the breadth of the sensing reach rather than to claim near-term operational impact.

6 Methane Emissions Monitoring: An End-to-End Case Study

This section assembles the full optimization–inference stack for single-photon LiDAR methane monitoring, and provides the paper’s running classical-versus-quantum comparison.

6.1 Hardware and Physical Principle

A SPAD detector registers individual photon-arrival (“click”) events backscattered from a methane plume; as noted below, a conventional SPAD reports threshold detection events via time-correlated single-photon counting (TCSPC) rather than resolving photon number. The DIAL principle uses two laser wavelengths: λ_{on} (strongly absorbed by CH_4) and λ_{off} (negligible absorption). The ratio of return photon counts encodes the column-integrated concentration via an idealized Beer–Lambert law, which in practice requires correction for wavelength-dependent system response, target reflectivity, atmospheric transmission, background light, and range effects:

$$\ln \frac{N_{\text{off}}}{N_{\text{on}}} = 2\sigma_{\text{CH}_4}(\lambda_{\text{on}}) \int_0^R c(r) dr, \quad (21)$$

where $\sigma_{\text{CH}_4}(\lambda_{\text{on}})$ is the absorption cross-section, R is the range, and $c(r)$ is the volumetric methane concentration along the beam path. Each pixel in the scan is a Poisson-distributed photon count—this is the quantum origin of the classical data, and the Poisson statistics set the fundamental shot-noise floor.

6.2 The Quantum Measurement Layer

In the low-photon regime relevant to eye-safe standoff sensing, photon statistics and detector architecture directly determine the Fisher information available for estimating absorption. Non-classical probe states and photon-number-resolving measurements can, under appropriate resource constraints and loss models, improve estimation precision relative to classical coherent-state strategies. The magnitude of any such quantum advantage depends critically on what resource is held fixed—incident photon number, absorbed photons, transmitted energy, acquisition time, or successful probe preparations—and on detector efficiency and loss; a naive comparison of a single-photon Fock state to a coherent state with the same mean photon number, for instance, can overstate the advantage, since the two states do not draw on the same photon budget once state preparation is accounted for. We therefore use the QFI/QCRB as a technology-independent benchmark rather than assigning a universal quantum-advantage factor.

Deployed-hardware caveat. Deployed SPADs typically provide threshold (“click”) rather than ideal photon-number-resolving detection. The resulting gap between the QFI and the classical Fisher information achieved by the deployed detector—quantified by Optimization 2’s measurement-loss decomposition—provides a natural measure of measurement loss and an engineering target for detector improvement.

6.3 Stage 1: Poisson MLE for Range and Absorption

Given photon arrival times from the TCSPC histogram, the range-resolved absorption profile is estimated via (18). The QCRB from Optimization 2 provides a variance floor: no estimator can achieve $\text{Var}(\hat{c}) < 1/(MF_Q^*)$ for a given photon budget M , playing the same role here as the shot count M elsewhere in the paper (Table 6). Here M denotes the effective number of independent information-bearing detection events under the adopted likelihood model; it need not equal the raw photon count when detector dead time, pile-up, background, or correlations across the shared latent absorption parameter are material. This bound separates hardware limitations from algorithmic inefficiency.

6.4 Stage 2: Tomographic Reconstruction

The ill-posed linear system $C = Kc$ is solved via the TV-regularized program (19). The condition number of the ray-integration matrix K depends critically on scan geometry. Minimizing $\text{cond}(K)$ over possible drone trajectories is a sensor placement subproblem with natural stochastic programming structure [Chaloner and Verdinelli, 1995, Birge and Louveaux, 2011]: the wind field u is uncertain, so the effective forward operator $\mathcal{K}(u) = K A(u)$ of Eq. (20) is stochastic:

$$\text{choose trajectory } \arg \min_{\text{traj}} \mathbb{E}_u[\text{cond}(\mathcal{K}(u))]. \quad (22)$$

6.5 Stage 3: Source Attribution via Advection-Diffusion Inversion

The volumetric methane field $c(\mathbf{x}, t)$ satisfies the advection-diffusion PDE:

$$\frac{\partial c}{\partial t} + \mathbf{u} \cdot \nabla c = D \nabla^2 c + S(\mathbf{x}), \quad (23)$$

where $\mathbf{u}$ is the wind field, D is turbulent diffusivity, and $S(\mathbf{x}) \geq 0$ is the unknown source distribution. Discretizing and inverting for S under wind uncertainty yields the robust minimax program (20) [Ben-Tal et al., 2009], in which S denotes the discretization of $S(\mathbf{x})$ on J candidate cells.

6.6 Stage 4: Adaptive Scan Allocation (Bandit Layer)

The drone’s scan-angle allocation is an adaptive bandit: at each time step, choose the next scan direction to maximize expected reduction in uncertainty about $S(\mathbf{x})$. Thompson sampling over a GP prior on the concentration field is the natural baseline; the Pareto-optimal balance between exploration and exploitation follows Simchi-Levi and Wang [2024].

An RL agent can additionally optimize the drone’s flight path (continuous action space) to jointly minimize reconstruction uncertainty and flight time—a multi-objective MDP with reward $-\text{Var}(S) - \kappa \cdot \text{flight cost}$ [Powell, 2022].

6.7 Stage 5: Operational Dispatch Decision

The estimated flux field $\hat{S}(\mathbf{x})$ with uncertainty quantification drives a repair crew dispatch optimization [Birge and Louveaux, 2011]:

$$\min_{\text{schedule}} \mathbb{E}_{S|\mathcal{D}} \left[\sum_k w_k \cdot \mathbf{1}[\text{leak } k \text{ unrepaired}] \cdot S_k \right] \quad \text{s.t. crew and time constraints}, \quad (24)$$

Table 5: The methane LiDAR end-to-end stack. Each row is one layer of the optimization–inference pipeline; the rightmost column shows how the quantum bound enters each stage.

Layer	Physics/ Engineering	OR/ML Method	Bounds Role
Probe design	Incident photon-number state $ n\rangle$; pulse energy	Optimization 1 / channel-QFI bound	$F_{Q,\text{noise}}^*$ sets variance floor
Measurement design	SPAD / PNR detector; detection POVM	Optimization 2 bounds achievable POVM performance	$F_C \leq F_Q$; measurement loss vs. ideal detection
Absorption extraction	Beer–Lambert; TC-SPC histogram	Poisson MLE (18)	QCRB bounds $\text{Var}(\hat{c})$
3-D reconstruction	Ray-integration $C = Kc$	TV regularization (19)	QCRB floor for absorption via (18)
Source attribution	Advection-diffusion PDE (23)	Robust minimax (20)	Absorption uncertainty (QCRB) propagated through wind-robust minimax
Scan allocation	Drone flight path	GP bandit / RL MDP (9)	Optimization 1 bound caps per-step IG via (12)
Dispatch	Crew scheduling	Stochastic integer program (24)	Flux UQ from Stage 3

where w_k weights leaks by global warming potential and S_k is the (random) flux at leak k , distributed according to its posterior given the data $\mathcal{D}$ propagated from Stage 3; $\hat{S}(\mathbf{x})$ denotes the corresponding posterior mean field. This is a stochastic integer program consuming the pipeline output [Shapiro et al., 2009].

6.8 Summary: The Full Stack

Table 5 summarizes the complete methane monitoring pipeline.

7 Benchmarking and Evaluation

The sensing-design optimizations and their SDP-computable bounds provide a principled upper bound on achievable precision—separating hardware limitations from algorithmic inefficiency, diagnosing where performance is being lost (in probe state preparation, in measurement choice, or irreversibly in decoherence), and guiding investment decisions. Fisher-information loss relative to the noiseless ideal decomposes additively across a four-level hierarchy:

$$F_Q^0 \geq F_{Q,\text{noise}}^* \geq F_Q(\rho_0) \geq F_C, \quad (25)$$

where F_Q^0 is the noiseless QFI, $F_{Q,\text{noise}}^*$ is the optimal QFI attainable under the noisy channel (the Optimization 1 value of Eq. (6); this is the same quantity written as F_Q^* elsewhere in the paper, spelled out here to distinguish it from the other three terms in the hierarchy), $F_Q(\rho_0)$ is the QFI actually achieved by the state ρ_0 that was prepared (which need not equal ρ_0^* ; $F_Q(\rho_0)$ is shorthand for $F_Q[\mathcal{E}_\phi(\rho_0)]$, the QFI of the channel output, since ρ_0 itself carries no information about ϕ), and

F_C is the classical Fisher information actually extracted by the deployed measurement. Because $F_{Q,\text{noise}}^*$ already reflects the noisy channel, the three losses are cleanly separable and additive:

$$\begin{aligned}
L_{\text{decoh}} &= F_Q^0 - F_{Q,\text{noise}}^* && \text{(irreducible; not recoverable by protocol design)} \\
L_{\text{prep}} &= F_{Q,\text{noise}}^* - F_Q(\rho_0) && \text{(closeable by Optimization 1)} \\
L_{\text{meas}} &= F_Q(\rho_0) - F_C && \text{(closeable by Optimization 2)} \\
F_Q^0 - F_C &= L_{\text{decoh}} + L_{\text{prep}} + L_{\text{meas}}. && (26)
\end{aligned}$$

These three losses connect to the RL optimality gap (17) and the point-estimator efficiency ratio (15) by the inequality chain:

$$\underbrace{\frac{1}{MF_{Q,\text{noise}}^*}}_{\text{Optimization 1 bound}} \leq \underbrace{\frac{1}{MF_Q(\rho_0)}}_{\text{actual prepared state}} \leq \underbrace{\frac{1}{MF_C}}_{\text{actual measurement}} \leq \text{Var}(\hat{\phi}). \quad (27)$$

NV-center example. For a single NV-center qubit probing field B via $\hat{H}_s = (\gamma_e B/2)\sigma_z$ under pure dephasing at rate Γ :

$$F_Q^0 = (\gamma_e \tau)^2, \quad F_{Q,\text{noise}}^* = F_Q(\rho_0) = F_C = (\gamma_e \tau)^2 e^{-2\Gamma\tau}.$$

Direct evaluation of the probe-design optimization (Optimization 1) shows that the Hadamard-prepared Ramsey state $(|0\rangle + |1\rangle)/\sqrt{2}$ is optimal under pure dephasing ($L_{\text{prep}} = 0$). The SLD-eigenbasis argument of Optimization 2 shows that a readout rotation followed by σ_z projective measurement is already optimal ($L_{\text{meas}} = 0$). The net fractional Fisher-information loss $L_{\text{decoh}}/F_Q^0 = 1 - e^{-2\Gamma\tau}$ is entirely decoherence loss L_{decoh} —irreducible given this noise model. The RL optimality gap $\delta_{\text{RL}} = 0$ when the Ramsey protocol is used; for more complex noise spectra, $\delta_{\text{RL}} > 0$ and RL provides genuine improvement [Krenn et al., 2023].

8 Research Opportunities for the INFORMS Data Science Community

8.1 The Classical–Quantum Interface

Every quantum sensor ultimately produces classical data that flows into a data science pipeline: feature extraction (absorption profiles, range histograms), model fitting (Bayesian posteriors, source distributions), uncertainty quantification, and downstream decisions. The quantum layer is characterized by probe- and measurement-design optimizations, often supported by SDP-computable bounds (Section 2). What quantum sensing offers the OR data scientist is unusually clean problem anatomy: exact forward models given by the Lindblad equation, SDP-certified precision ceilings, and a direct correspondence between algorithmic reward and physical information gain.

The most important classical–quantum interface points are:

1. **Measurement selection**—Optimization 2 determines which POVM to apply, a pre-measurement design decision.
2. **Parameter estimation**—maximum likelihood for static stimuli, Bayesian bandit updating (Section 4) for adaptive protocols.
3. **Uncertainty quantification**—monitoring the gap between achieved variance and the QCRB in real time diagnoses hardware degradation and suboptimal protocol choices.

4. **Downstream decision**—for methane monitoring, the estimated concentration field drives repair crew dispatch (24); for MEG, it drives surgical planning. These are stochastic optimization problems consuming the pipeline output [Shapiro et al., 2009, Birge and Louveaux, 2011].

8.2 Better Algorithms for the Sensing Optimizations

Where the probe-design and measurement-design problems of Section 2 admit exact or bounding SDP/conic-program formulations, calling a generic interior-point solver is far from the end of the story. The proof of concept is Peña et al. [2025], who demonstrated for quantum entanglement SDPs that tailored methods achieve dramatically better scaling than MOSEK applied off-the-shelf. Related structural properties—Schur-complement block constraints, Lyapunov-type structure, and low-rank optimal solutions—recur across many SDP formulations of sensing problems and are a natural target for the same kind of tailoring.

Exploiting Lyapunov–Schur structure. SDP formulations of the probe-design problem often introduce auxiliary matrices via a Schur complement, producing a block-PSD constraint with Lyapunov-equation structure. For particular Lyapunov–Schur formulations, exploiting structure can substantially reduce the dominant linear-algebra cost relative to a dense generic SDP implementation. The attainable scaling is formulation-, sparsity-, and accuracy-dependent; deriving and testing these reductions for sensing-specific SDPs is an open research problem.

Symmetry reduction. When the sensing Hamiltonian has a known symmetry group, ρ_0 block-diagonalizes under the group action. For the permutation group $\mathcal{S}_N$ over N identical probe qubits, the resulting SDP grows polynomially in N rather than exponentially [Tavakoli et al., 2024]. Combining symmetry reduction with the tailored methods of Peña et al. [2025] is an open problem.

Warm-starting in adaptive protocols. In an adaptive sensing protocol, Optimization 1 is solved once per measurement stage with a small update to the noise model. The solution at stage t is an excellent warm start for stage $t + 1$. The low-rank structure of typical probe states suggests that active-set or rank-revealing factorization approaches may be effective—an open problem in computational optimization.

Mixed-integer SDPs for error-corrected sensing. When quantum error correction (QEC) overheads are incorporated, Optimization 1 must include them as additional constraints, creating mixed-integer semidefinite programs that combine code selection (discrete) with probe state optimization (continuous PSD)—a new class of hard problems.

8.3 Non-SDP Problems

Bandit algorithms for adaptive measurement. Open problems include: (i) contextual bandits when the noise model $\mathcal{E}$ changes with environmental conditions; (ii) batched bandits when measurements are taken in parallel across N probe elements, addressing the Pareto trade-off of Simchi-Levi and Wang [2024]; (iii) best-arm identification under the per-shot and mission-time constraints of (9).

RL for pulse-sequence design. Open problems include: (i) sample efficiency—current RL agents require many simulated episodes; (ii) robustness to model mismatch, with off-policy evaluation methods [Wang et al., 2024] providing rigorous guarantees; (iii) multi-objective sequences that simultaneously maximize T_2 , minimize pulse power, and satisfy hardware constraints.

ML surrogates for real-time posterior inference. The broader opportunity is simulation-based inference (SBI) [Cranmer et al., 2020]: given a forward model of the quantum sensor, train a normalizing flow to approximate the posterior directly. The QFI bound provides a principled local efficiency benchmark for a specified point estimator derived from the surrogate, via Eq. (15), while posterior quality should be assessed separately through calibration, coverage, proper scoring rules, and posterior predictive checks.

Stochastic optimization for networked sensors. When multiple quantum sensors are deployed in a network—LiDAR platforms across a gas field, gravimeters in a geophysical survey, OPM arrays on a patient’s scalp—the joint sensing problem is a multi-source data assimilation problem. Key OR subproblems include online allocation of sensing resources [Shapiro et al., 2009] and robust optimization for meteorological model mismatch [Ben-Tal et al., 2009].

Broader governance implications. As quantum-limited sensors approach fundamental detection floors in high-stakes absence-of-signal settings—such as cardiac monitoring in donation after circulatory death protocols—the measurement chain’s design acquires legal and clinical consequence [Roberts et al., 2026]. This creates an OR research agenda: designing reference detection chains as constrained optimization problems, specifying evidentiary rules as decision rules under uncertainty [Gan et al., 2025], and bounding adversarial reanalysis risk via prespecified hypothesis testing with controlled error rates—all standard INFORMS data science toolkit items.

LLMs for quantum sensing workflows. At the human–machine interface, LLMs offer a practical near-term role: translating high-level operational objectives into sensing protocol parameters, and interpreting diagnostic output from SDPs and bandit algorithms in plain language for field operators.

9 Conclusion

Quantum sensing has matured from a laboratory specialty to a field with deployed commercial and governmental applications. Atomic clocks underpin GPS. SQUID magnetometers inform epilepsy surgery planning. Single-photon LiDAR systems monitor methane emissions from gas infrastructure, and the optimization–inference stack of this paper traces that application end to end, from photon statistics through SDP-computable bounds to operational crew dispatch.

This paper has had two purposes, and it is best read as a position paper rather than a technical SDP contribution. As a primer for the INFORMS data science community, it showed that quantum sensing is not a peripheral physics specialty but a source of concrete optimization and inference problems with unusually clean structure: exact forward models given by the Lindblad equation, SDP-certified precision ceilings, and a direct correspondence between algorithmic reward and physical information gain. The two sensing-design optimizations of Section 2 serve as enabling design tools—given a noise model, Optimization 1 returns the probe state to prepare and Optimization 2 returns the measurement to implement—rather than as the paper’s primary novelty.

As a research agenda, it mapped where the INFORMS data science community can go beyond formulating problems and calling solvers. On the SDP side, the Lyapunov–Schur structure common to SDP formulations of the probe-design problem, the permutation symmetries of multi-qubit arrays, and the warm-start structure of adaptive protocols are all unexploited by current software; Peña et al. [2025] demonstrated that tailored methods can achieve dramatic improvements for structurally related quantum SDPs. On the non-SDP side, adaptive measurement admits a bandit formulation whose local-Gaussian information-gain approximation is capped by QFI as in Eq. (12); pulse sequence design is an RL problem with demonstrated results [Krenn et al., 2023] and a formal MDP specification; MEG source imaging and LiDAR plume reconstruction call for ML surrogates and simulation-based inference; and networked sensor data assimilation is a multi-agent stochastic optimization problem.

What makes quantum sensing an unusually attractive domain for the INFORMS data science community is not that it requires exotic new methods, but that it offers a uniquely well-specified arena for applying methods the community already knows: the physics provides exact forward models, the SDPs provide certified precision ceilings, and the gap decomposition tells the data scientist exactly where to invest algorithmic effort.

Quantum sensing is therefore not merely a hardware story. It is an optimization story constrained by physics, an inference story constrained by information, and increasingly, an operational data science story constrained by time, uncertainty, and scale.

A Notation

Table 6 collects the principal symbols used throughout the paper. Note: $\text{Tr}(\cdot)$ denotes the matrix trace (sum of diagonal entries) throughout.

B Ramsey Interferometry, Quantum Limits, and the Cramér–Rao Gap

B.1 Ramsey Interferometry: Full Derivation

Ramsey interferometry is the canonical single-qubit sensing protocol and provides the calibration point for the SDP formulations. Setup: take a constant longitudinal stimulus $V_\perp = 0$, $V_\parallel = V_0$, and use the Hadamard gate $\hat{U}_H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ for both initialization and readout.

Initialization. The prepared state is $|\psi_{\text{init}}\rangle = \hat{U}_H|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$.

As in Section 2.3, γ here is expressed in angular-frequency units, so the signal term γV_0 carries no explicit factor of $\hbar$; ΔE remains an energy and enters as $\Delta E/\hbar$, matching Eq. (3).

Time evolution. Under the signal Hamiltonian with $\Delta E \rightarrow \Delta E + \gamma V_0$: $|\psi(t)\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\theta}|1\rangle)$, where $\theta = t(\Delta E/\hbar + \gamma V_0)$.

Readout and transition probability. Applying $\hat{U}_H$ again and measuring in the computational basis gives:

$$p_{0 \rightarrow 1} = \frac{1}{2} \left[1 - \cos t(\Delta E/\hbar + \gamma V_0) \right]. \quad (28)$$

Fisher information. For this Bernoulli measurement, the classical Fisher information (CFI) computed from the derivative with respect to V_0 at the operating point $\theta = \pi/2$ (where $\sin \theta = 1$

Table 6: Principal notation.

Symbol	Meaning
$\text{Tr}(\cdot)$	Matrix trace (sum of diagonal entries)
$ \psi\rangle, \langle\psi $	Quantum state ket and conjugate transpose
$\rho \succeq 0, \text{Tr}(\rho) = 1$	Density operator
$F_Q(\rho_\phi)$	Quantum Fisher information
$F_C(p_\phi)$	Classical Fisher information
L	Symmetric logarithmic derivative (SLD)
$\mathcal{E}$	Parameter-independent noise (decoherence) map (CPTP)
$\mathcal{E}_\phi$	Parameter-dependent sensing channel (encoding by $H_s = \phi G$ plus noise); $\rho_\phi = \mathcal{E}_\phi(\rho_0)$
$\{E_i\}$	POVM: $E_i \succeq 0, \sum_i E_i = I$
G	Generator; $H_s = \phi G$
τ, T_2	Sensing time; coherence time
N, M	Probe particles; measurement repetitions (shots), or, in Section 6, effective independent detection events (see discussion there)
$\widehat{\Delta\phi}_{\text{SQL}}$	Standard quantum limit
$\widehat{\Delta\phi}_{\text{HL}}$	Heisenberg limit
$\phi; \phi \in \mathbb{R}^d$	Scalar parameter; vector of d parameters estimated jointly
Σ	Estimator covariance matrix, induced by the POVM and estimator (not a free decision variable)
W	Parameter weight matrix
δ_{RL}	RL optimality gap, Eq. (17)
$\bar{n}$	Mean photon number per pulse
r_i, n_i, μ	Range of TCSPC bin i ; photon count in bin i ; intensity function (expected counts per bin), Eq. (18)
c, C, K	Voxel concentrations; path-integrated measurements; ray-integration matrix, Eq. (19)
$S, \mathcal{K}(u), \mathcal{U}$	Source emission rates; wind-dependent forward operator $K A(u)$; wind uncertainty ellipsoid, Eq. (20)
$\lambda_1, \lambda_2, \lambda_3$	Regularization parameters in Eqs. (18)–(20)

and $p_{0 \rightarrow 1}(1 - p_{0 \rightarrow 1}) = 1/4$ is:

$$F_C = \frac{(\partial_{V_0} p_{0 \rightarrow 1})^2}{p_{0 \rightarrow 1}(1 - p_{0 \rightarrow 1})} = (t\gamma)^2. \quad (29)$$

This equals the QFI F_Q for the Hadamard-prepared state, confirming that the Ramsey protocol is *optimal* for a noiseless longitudinal stimulus: no alternative initialization or measurement can do better.

B.2 Shot Noise, the Heisenberg Limit, and Entanglement

With N independent probes and M shots, additivity of Fisher information across independent probes together with the classical Cramér–Rao bound gives the *standard quantum limit* (SQL):

$$\widehat{\Delta\phi}_{\text{SQL}} = \frac{1}{\sqrt{NM}}. \quad (30)$$

When all N probes are entangled in the Greenberger–Horne–Zeilinger (GHZ) state $|\Psi_{\text{GHZ}}\rangle = \frac{1}{\sqrt{2}}(|0\rangle^{\otimes N} + |1\rangle^{\otimes N})$, the joint system accumulates phase $N\phi$ per cycle, yielding the *Heisenberg limit* (HL):

$$\widehat{\Delta\phi}_{\text{HL}} = \frac{1}{N\sqrt{M}}, \quad (31)$$

a quadratic reduction in the required number of probes for fixed target precision (equivalently, a $\sqrt{N}$ reduction in standard deviation at fixed N) [Giovannetti et al., 2011]. The GHZ state is inseparable (Proposition 1); this inseparability is what enables the gain.

Proposition 1 (Inseparability of the $N = 2$ GHZ state [Nielsen and Chuang, 2010]). *The state $(|00\rangle + |11\rangle)/\sqrt{2}$ cannot be written as $|\psi_1\rangle \otimes |\psi_2\rangle$ for any $|\psi_k\rangle \in \mathbb{C}^2$.*

Proof. Writing $|\psi_k\rangle = a_k|0\rangle + b_k|1\rangle$, matching coefficients to $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ requires $a_1a_2 = b_1b_2 = 1/\sqrt{2}$ and $a_1b_2 = b_1a_2 = 0$. The zero conditions force $a_1a_2 = 0$ or $b_1b_2 = 0$, a contradiction. $\square \square$

B.3 The Sensing Protocol and the Cramér–Rao Gap

After free evolution under $H_s = \phi G$ for time τ , the probe state is $|\psi_\phi\rangle = e^{-iG\phi\tau}|\psi_0\rangle$. A readout pulse U_R transforms the basis, and a computational-basis measurement yields outcome 0 with probability $p_0(\phi) = |\langle 0|U_R|\psi_\phi\rangle|^2$. From M independent repetitions, the MLE satisfies $\text{Var}(\hat{\phi}) \geq 1/(MF_C)$ with $F_C = [p'_0(\phi)]^2/[p_0(\phi)(1 - p_0(\phi))]$.

For the Ramsey example, $F_C = (\gamma_e\tau)^2 = F_Q$, confirming optimality. The three-component gap decomposition of Section 7 follows from the inequality chain (27):

$$\underbrace{\frac{1}{MF_{Q,\text{noise}}^*}}_{\text{Optimization 1 bound}} \leq \underbrace{\frac{1}{MF_Q(\rho_0)}}_{\text{actual prepared state}} \leq \underbrace{\frac{1}{MF_C}}_{\text{actual meas.}} \leq \text{Var}(\hat{\phi}).$$

For the NV center under pure dephasing, the analytic solutions of Optimization 1 and Optimization 2 both give zero recoverable loss; the entire gap is decoherence loss $1 - e^{-2\Gamma\tau}$, irreducible given the noise model.

C Hardware Landscape

Quantum sensing spans hardware maturity from atomic clocks at TRL 9 to NV-center nanoscale imagers at TRL 5. Public programs in the United States, United Kingdom, European Union, and other countries support quantum-sensing R&D, including the US DARPA Transformative Quantum Sensing program, UK National Quantum Technologies Programme, EU Quantum Flagship, and national programs in France, Germany, Australia, Japan, and Singapore. Commercial firms include AOSense, Exail (formerly Muquans), and GEM Systems (gravimetry, atomic clocks, magnetometry); QLM Technology and Bridger Photonics (LiDAR gas detection); and Q-CTRL (quantum control software).

D Glossary of Acronyms

Table 7 lists acronyms from quantum sensing literature for a quick reference.

Table 7: Quantum sensing and quantum information acronyms.

Acronym	Expansion
AQG	Absolute Quantum Gravimeter
CPMG	Carr–Purcell–Meiboom–Gill (dynamical decoupling pulse sequence)
CPTP	Completely Positive, Trace-Preserving (map)
DIAL	Differential Absorption LiDAR
GHZ	Greenberger–Horne–Zeilinger (entangled state)
HL	Heisenberg Limit
LiDAR	Light Detection and Ranging
LIGO	Laser Interferometer Gravitational-Wave Observatory
MAB	Multi-Armed Bandit
MDP	Markov Decision Process
MEG	Magnetoencephalography
NV	Nitrogen-Vacancy (center in diamond)
OPM	Optically Pumped Magnetometer
POVM	Positive Operator-Valued Measure
Q-INS	Quantum Inertial Navigation System
QCRB	Quantum Cramér–Rao Bound
QEC	Quantum Error Correction
QFI	Quantum Fisher Information
RL	Reinforcement Learning
SBI	Simulation-Based Inference
SLD	Symmetric Logarithmic Derivative
SPAD	Single-Photon Avalanche Diode
SQL	Standard Quantum Limit
SQUID	Superconducting Quantum Interference Device
TCSPC	Time-Correlated Single-Photon Counting
TRL	Technology Readiness Level

References

- Barry, J. F., Schloss, J. M., Bauch, E., Turner, M. J., Hart, C. A., Pham, L. M., and Walsworth, R. L. (2020). Sensitivity optimization for NV-diamond magnetometry. *Reviews of Modern Physics*, 92(1):015004.
- Ben-Tal, A., El Ghaoui, L., and Nemirovski, A. (2009). *Robust Optimization*. Princeton University Press.
- Birge, J. R. and Louveaux, F. (2011). *Introduction to Stochastic Programming*. Springer, 2nd edition.
- Bongs, K. et al. (2019). Taking atom interferometric quantum sensors from the laboratory to real-world applications. *Nature Reviews Physics*, 1(12):731–739.
- Boto, E., Holmes, N., Leggett, J., Roberts, G., Shah, V., Meyer, S. S., Muñoz, L. D., Mullinger, K. J., Tierney, T. M., Bestmann, S., Barnes, G. R., Bowtell, R., and Brookes, M. J. (2018). Moving magnetoencephalography towards real-world applications with a wearable system. *Nature*, 555(7698):657–661.
- Chaloner, K. and Verdinelli, I. (1995). Bayesian experimental design: A review. *Statistical Science*, 10(3):273–304.

- Cranmer, K., Brehmer, J., and Louppe, G. (2020). The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062.
- Degen, C. L., Reinhard, F., and Cappellaro, P. (2017). Quantum sensing. *Reviews of Modern Physics*, 89(3):035002.
- Demkowicz-Dobrzański, R., Kołodyński, J., and Guţă, M. (2012). The elusive Heisenberg limit in quantum-enhanced metrology. *Nature Communications*, 3:1063.
- Gan, K., Jia, S., Li, A., and Tayur, S. (2025). Towards a liquid biopsy: Greedy approximation algorithms for active sequential hypothesis testing. *Management Science*.
- Giovannetti, V., Lloyd, S., and Maccone, L. (2011). Advances in quantum metrology. *Nature Photonics*, 5:222–229.
- Hauke, P., Heyl, M., Tagliacozzo, L., and Zoller, P. (2016). Measuring multipartite entanglement through dynamic susceptibilities. *Nature Physics*, 12:778–782.
- Hayashi, M. and Ouyang, Y. (2023). Tight Cramér–Rao type bounds for multiparameter quantum metrology through conic programming. *Quantum*, 7:1094.
- Krenn, M., Landgraf, J., Foesel, T., and Marquardt, F. (2023). Artificial intelligence and machine learning for quantum technologies. *Physical Review A*, 107:010101.
- Liu, J., Chen, Z., and Yuan, H. (2024). Fully-optimized quantum metrology: Framework, tools, and applications. *Advanced Quantum Technologies*, 7:2400094.
- Mathew, G., Silva, S. L. L., Jain, A., Mohan, A., Adroja, D. T., Sakai, V. G., Tomy, C. V., Banerjee, A., Goreti, R., Aswathi, V. N., Singh, R., and Jaiswal-Nagar, D. (2020). Experimental realization of multipartite entanglement via quantum Fisher information in a uniform antiferromagnetic quantum spin chain. *Physical Review Research*, 2:043329.
- MIT News Office (2025). LiDAR helps gas industry find methane leaks and avoid costly losses. Published September 12, 2025. <https://news.mit.edu/2025/lidar-helps-gas-industry-find-methane-leaks-avoid-costly-losses-0912>.
- Muquans/Exail (2021). Absolute quantum gravimeter (AQG). Technical report, Exail. <https://www.exail.com/>.
- Nielsen, M. A. and Chuang, I. L. (2010). *Quantum Computation and Quantum Information*. Cambridge University Press, 10th anniversary edition.
- National Institute of Standards and Technology (2025). Atomic time for a networked world. <https://www.nist.gov/atomic-clocks/atomic-time-networked-world>.
- Peña, J., Siddhu, V., and Tayur, S. (2025). Tailored first-order and interior-point methods and a new semidefinite programming hierarchy for entanglement detection. arXiv preprint arXiv:2508.05854.
- Powell, W. B. (2022). *Reinforcement Learning and Stochastic Optimization: A Unified Framework for Sequential Decisions*. John Wiley & Sons.
- QLM Technology Ltd (2025). Quantum LiDAR for methane emissions monitoring: EPA alternative test method qualification. <https://qlmtec.com/>.

- Roberts, J., Tayur, S., and Kim, T. W. (2026). The next quantum revolution will challenge the evidentiary basis of the dead donor rule. Working paper, Carnegie Mellon University and University of California, San Francisco.
- Russo, D., Van Roy, B., Kazerouni, A., Osband, I., and Wen, Z. (2018). A tutorial on Thompson sampling. *Foundations and Trends in Machine Learning*, 11(1):1–96.
- Scheie, A., Laurell, P., Samarakoon, A. M., Lake, B., Nagler, S. E., Granroth, G. E., Okamoto, S., Alvarez, G., and Tennant, D. A. (2021). Witnessing entanglement in quantum magnets using neutron scattering. *Physical Review B*, 103:224434.
- Shapiro, A., Dentcheva, D., and Ruszczyński, A. (2009). *Lectures on Stochastic Programming: Modeling and Theory*. SIAM/MPS.
- Shen, Z., Ding, S., Zhao, Z., Evangelista, F. A., and Wang, Y. (2026). Witnessing spin-orbital entanglement using resonant inelastic x-ray scattering. *Physical Review Letters*, 136:116504.
- Shi, H.-L., Guan, X.-W., and Yang, J. (2024). Universal shot-noise limit for quantum metrology with local Hamiltonians. *Physical Review Letters*, 132(10):100803.
- Siddhu, V. and Tayur, S. (2022). Five starter pieces: Quantum information science via semidefinite programs. *Tutorials in Operations Research*, pages 59–92. INFORMS.
- Simchi-Levi, D. and Wang, C. (2024). Multi-armed bandit experimental design: Online decision-making and adaptive inference. *Management Science*. doi:10.1287/mnsc.2023.00492.
- Tang, Y., Li, A., Scheller-Wolf, A., and Tayur, S. (2025a). Multi-armed bandits with endogenous learning curves: An application to split liver transplantation. *Manufacturing & Service Operations Management*.
- Tang, Y., Scheller-Wolf, A., Tayur, S., Perito, E., and Roberts, J. (2025b). Split liver transplantation: An analytical decision support model. *Operations Research*.
- Tavakoli, A., Pozas-Kerstjens, A., Brown, P., and Araújo, M. (2024). Semidefinite programming relaxations for quantum correlations. *Reviews of Modern Physics*, 96(4):045006.
- Wang, J., Gao, R., and Zha, H. (2024). Reliable off-policy evaluation for reinforcement learning. *Operations Research*, 72(2):699–716.
- Yu, M., Liu, Y., Yang, P., Gong, M., Cao, Q., Zhang, S., Liu, H., Heyl, M., Ozawa, T., Goldman, N., and Cai, J. (2022). Quantum Fisher information measurement and verification of the quantum Cramér–Rao bound in a solid-state qubit. *npj Quantum Information*, 8:56.