

Quantum Sensing as an Optimization–Inference Stack

Sridhar Tayur

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

June 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; superconducting magnetometers guide epilepsy surgery; single-photon LiDAR systems detect methane leaks at parts-per-million from 200 meters. This position paper argues that quantum sensing is naturally viewed as an optimization–inference stack whose central design problems are semidefinite programs (SDPs) and whose downstream layers connect directly to statistical inference and operational decision-making familiar to the INFORMS data science community—adaptive bandits, reinforcement learning, inverse problems, stochastic programming, and ML surrogates. The two SDP formulations are developed from first principles and shown to bound every other quantity in the pipeline; they are best understood as oracle tools that certify precision ceilings rather than as the paper’s primary novelty. An end-to-end methane emissions case study traces the complete stack from photon physics through SDP design to operational crew dispatch, providing a concrete 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, detecting methane plumes at parts-per-million (ppm) concentration from 200 meters and streaming a three-dimensional map of emission rates to a cloud processing pipeline that 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 **SDP layer** characterizes—before any experiment is conducted—the best probe state and measurement that physics permits for a given noise model. The resulting quantum Cramér–Rao bound (QCRB) is a ceiling on estimation precision 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 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]. 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 SDP formulations, 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 path to field use not yet established.

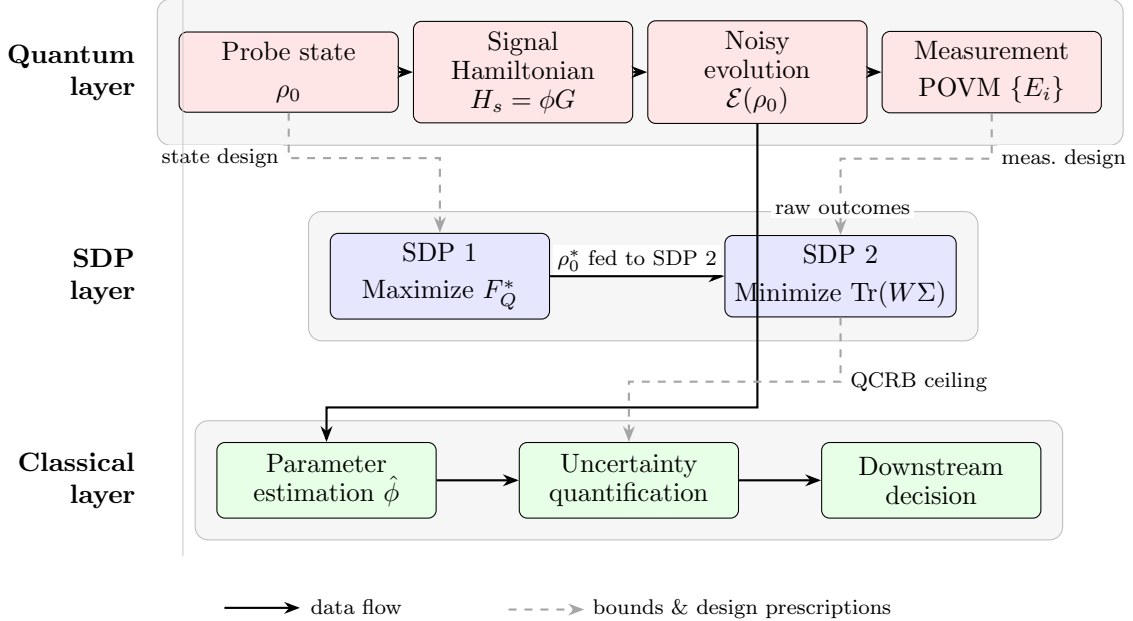


Figure 1: The quantum sensing pipeline as an optimization–inference stack. **Quantum layer** (top): the probe state ρ_0 is prepared, exposed to the signal Hamiltonian $H_s = \phi G$, undergoes noisy evolution $\mathcal{E}(\rho_0)$, and is measured by POVM $\{E_i\}$. **SDP layer** (center): SDP 1 optimizes ρ_0 to maximize the quantum Fisher information F_Q^* ; SDP 2 uses the resulting ρ_0^* to optimize the measurement POVM, minimizing the weighted estimation error $\text{Tr}(W\Sigma)$. Together they establish the QCRB that limits every quantity in the layer below. **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 SDP layer.

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, and the constraints that arise in sensing protocols are linear in ρ .

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 $p_i = \text{Tr}(\rho \Pi_i)$. Once probabilities are defined, the data collected by a quantum sensor is a classical random variable and standard statistical methods apply.

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]

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 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 SDP 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 SDP 1. After signal interaction and decoherence, the state becomes $\rho_{\phi} = \mathcal{E}(\rho_0)$; 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 noise map \mathcal{E} and the signal Hamiltonian.

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.

2.3 The Four-Step Sensing Protocol

1. **Initialize.** Apply \hat{U}_{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}_{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 is exactly $\hat{H}_s = (\gamma_e B/2)\sigma_z$, a direct instantiation of $H_s = \phi G$ with $\phi = \gamma_e B$ (the magnetic field to be measured) and $G = \sigma_z/2$. 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). Figure 2 at the end of this section maps the NV example to each layer of Figure 1.

2.4 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 improvement, 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 ($N = 1$), the SQL sets the per-shot floor; entanglement across an NV array can in principle reach the HL. Full derivations, including the inseparability proof for the GHZ state and the Ramsey Fisher-information calculation, are given in Appendix B.

2.5 Quantum Fisher Information and the Cramér–Rao Bound

The *quantum Fisher information* (QFI) $F_Q(\rho_\phi)$ is the upper limit on the classical Fisher information F_C optimized over all possible measurements. The *quantum Cramér–Rao bound* (QCRB) states that for any unbiased estimator 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, the QFI equals four times the variance of the generator: $F_Q = 4 \text{Var}_{\psi_0}(G)$. 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 SDP formulation of the next subsection becomes essential.

2.6 Why Semidefinite Programs Arise Inevitably

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]. The connection to quantum sensing is structural, not coincidental: four properties of quantum mechanics force the design problem into exactly this form. First, density matrices are PSD by definition: the space of valid probe states is a convex body inside the PSD cone, exactly the feasible set of an SDP. Second, physically realizable measurements—POVMs—are collections of PSD operators summing to identity; optimizing over measurements is again optimization over a PSD cone. Third, the QCRB (4) is a convex relaxation of the estimation variance, linear in the inverse of the QFI. Fourth, the key technical step that converts the nonlinear objective $\text{Tr}(\rho_\phi L^2)$ into a tractable form introduces an auxiliary Hermitian matrix X and enforces the block PSD constraint

$$\begin{pmatrix} \rho_\phi & X \\ X & \rho_\phi \end{pmatrix} \succeq 0,$$

which, by the Schur complement lemma, is equivalent to $\rho_\phi \succeq 0$ and $\rho_\phi - X\rho_\phi^{-1}X \succeq 0$. This is the same algebraic idea used to convert quadratic Lyapunov stability conditions into linear matrix inequalities in control theory; see Demkowicz-Dobrzański et al. [2012] for the full derivation.

2.7 SDP Formulation 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?

When the probe undergoes decoherence, $\rho_\phi = \mathcal{E}(\rho_0)$ is mixed, where \mathcal{E} is a completely positive, trace-preserving (CPTP) linear map. 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 $F_Q = \text{Tr}(\rho_\phi L^2)$.

SDP 1 (Optimal Probe State Under Decoherence) [Demkowicz-Dobrzański et al., 2012, Liu et al., 2024]

$$\begin{aligned} & \underset{\rho_0, X}{\text{maximize}} && \text{Tr}(X) \\ & \text{subject to} && \begin{pmatrix} \rho_\phi & X \\ X & \rho_\phi \end{pmatrix} \succeq 0 \\ & && \rho_\phi = \mathcal{E}(\rho_0), \quad \rho_0 \succeq 0, \quad \text{Tr}(\rho_0) = 1. \end{aligned}$$

The optimal value is $F_Q^*/4$; the optimal ρ_0^* is the probe state to prepare.

Decision variables—physical meaning.

- ρ_0 : the experimenter’s choice of probe preparation.
- $\rho_\phi = \mathcal{E}(\rho_0)$: the noisy probe state after interaction. The noise map \mathcal{E} is a *parameter*, not a decision variable.
- X : an auxiliary matrix introduced to linearize $\text{Tr}(\rho_\phi L^2)$.

The program has $O(d^2)$ real decision variables ($d = 2^n$). For $n = 1, 2, 3, 4$ qubits: 4, 16, 64, 256 variables—tractable on a laptop.

Duality and the precision bound. Every feasible dual point provides an upper bound on F_Q^* ; every feasible primal point provides a lower bound via a constructive probe state. Slater’s constraint qualification holds whenever \mathcal{E} maps the interior of the density operator cone to itself—satisfied by all physically realistic decoherence models—ensuring strong duality [Tavakoli et al., 2024].

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 $e^{-\Gamma\tau}$. SDP 1 confirms that the Hadamard-prepared state $(|0\rangle + |1\rangle)/\sqrt{2}$ is optimal: 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 treats SDP 1 as a black box 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. The formulation shown here is a simplified presentation of the SDP representation of QFI developed in Demkowicz-Dobrzański et al. [2012] and Liu et al. [2024]; interested readers are referred there for the full derivation.

2.8 SDP Formulation 2: Designing the Optimal Measurement

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

A positive operator-valued measure (POVM) is the most general description of a quantum measurement. 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.

For single-parameter estimation the optimal measurement is the projective measurement in the eigenbasis of the SLD operator L . For the multiparameter case—simultaneously estimating all three spatial components of a magnetic field, for example—no single POVM achieves all bounds simultaneously in general [Hayashi and Ouyang, 2023, Liu et al., 2024].

SDP 2 (Optimal POVM for Multiparameter Estimation) [Hayashi and Ouyang, 2023, Liu et al., 2024]

$$\begin{aligned} & \underset{\{E_i\}, \Sigma}{\text{minimize}} && \text{Tr}(W\Sigma) \\ & \text{subject to} && \Sigma_{jk} \geq \frac{1}{2} \text{Re} \text{Tr}[\rho_\phi(L_j E_i L_k + L_k E_i L_j)] \quad \forall i, j, k \\ & && E_i \succeq 0 \quad \forall i, \quad \sum_i E_i = I, \quad \Sigma \succeq 0. \end{aligned}$$

Decision variables—physical meaning.

- $\{E_i\}$: the POVM—which outcomes are possible and with what probabilities.
- Σ : the estimator covariance matrix—the error budget.
- $W \succ 0$: the application-specific weight matrix. For navigation, position errors may be weighted more heavily than orientation; for magnetoencephalography (MEG), the direction of maximum neuronal activity dominates.

SDPs 1 and 2 together form a two-stage design pipeline. The gap between any practical protocol’s performance and the bound decomposes into state-preparation loss (closeable by SDP 1) and measurement loss (closeable by SDP 2).

Take-home for non-quantum readers. A practitioner who treats SDP 2 as a black box receives the optimal measurement operators $\{E_i^*\}$ (an instruction for what to measure) and the minimum achievable covariance matrix Σ^* (the error budget). Together, SDP 1 and SDP 2 serve as oracle

Table 2: NV-center magnetometry running example mapped to the optimization–inference stack of Figure 1. Each row is one pipeline stage; the SDP column records what SDP 1 or SDP 2 certifies. 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	SDP layer	Classical layer
State design	Prepare $(0\rangle+ 1\rangle)/\sqrt{2}$ via Hadamard gate	SDP 1 confirms Hadamard state is optimal	State-preparation loss = 0
Signal encoding	Field B imprints phase $\gamma_e B \tau$ (Ramsey protocol)	$F_Q^* = (\gamma_e \tau)^2 e^{-2\Gamma\tau}$ (SDP 1 output)	QCRB floor: $\text{Var}(\hat{B}) \geq 1/(MF_Q^*)$
Measurement	σ_z projective readout via optical spin detection	SDP 2 confirms σ_z is the optimal POVM	Measurement loss = 0
Pulse design	Dynamical decoupling suppresses dephasing rate Γ	RL optimality gap benchmarked against F_Q^*	Bandit 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}

tools that certify 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 center, σ_z readout is the eigenbasis of the SLD under pure dephasing, so SDP 2 confirms that the standard projective measurement is already optimal: measurement loss is also zero. The full precision gap is therefore pure decoherence loss, quantified exactly by the inequality chain in Section 7.

3 Hardware Platforms

Having established the optimization–inference framework and its SDP-certified precision ceilings, we now survey the major sensing platforms that instantiate it, noting the specific algorithmic bottleneck each one presents.

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

Table 3: Quantum sensing hardware platforms. fT = femtotesla; pT = picotesla; ppm = parts per million; 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 at 200 m	No	6-8	[Field trial]

is Poisson-limited and ill-conditioned in the spatial reconstruction; its bottleneck is compressed sensing and source attribution under meteorological uncertainty.

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 exactly $\hat{H}_s = (\gamma_e B/2)\sigma_z$ —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.

3.2 Inertial and Gravitational Sensing

Atom interferometers [Field trial] accumulate differential phase under gravity or rotation. 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—the most mature quantum sensing platform.

Optical lattice clocks [Lab/prototype] reach fractional frequency uncertainties below 10^{-18} ; spin squeezing has demonstrated genuine operational improvement in stability [Nature Collections, 2025] but lattice clocks remain laboratory instruments.

3.4 Photonic Sensing: Single-Photon LiDAR

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

4 Sensing Algorithms and Signal Processing

The two SDPs above 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 SDP 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:

$$p(\phi | x_{1:t}) \propto p(x_t | \phi, \xi_t) p(\phi | x_{1:t-1}). \quad (5)$$

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_{\max} plays the role of the budget: once the qubit decoheres, no further useful measurements can be made in that experimental run.

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:

$$\max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=1}^M \text{IG}(\xi_t, x_t) \right] \quad \text{subject to} \quad \sum_{t=1}^M \tau(\xi_t) \leq T_{\max}, \quad (6)$$

where the per-step information gain is

$$\text{IG}(\xi_t, x_t) = H[p(\phi | x_{1:t-1})] - \mathbb{E}_{x_t} [H[p(\phi | x_{1:t})]], \quad (7)$$

$H[\cdot]$ denotes Shannon entropy, $\tau(\xi_t)$ is the duration of measurement ξ_t , and T_{\max} is the coherence-time budget.

The SDP layer enters as an upper bound on per-step gain. For a locally Gaussian posterior approximation with variance σ_t^2 , the information gain satisfies [Giovannetti et al., 2011, Russo et al.,

2018]:

$$\text{IG}(\xi_t, x_t) \lesssim \frac{1}{2} \log(1 + F_Q(\rho_\phi) \sigma_t^{-2}), \quad (8)$$

where F_Q is computed by SDP 1.

Thompson sampling [Russo et al., 2018] applies directly:

$$\tilde{\phi}_t \sim p(\phi | x_{1:t-1}), \quad \xi_t = \arg \max_{\xi \in \Xi} p(\tilde{\phi}_t | \xi). \quad (9)$$

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 (9). 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 (6). 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:

$$\hat{\theta} = \arg \min_{\theta} \mathbb{E}_{(\phi, x_{1:t}) \sim p} [D_{\text{KL}}(p(\phi | x_{1:t}) \| q_\theta(\phi | x_{1:t}))]. \quad (10)$$

The SDP bound on F_Q^* provides a ground-truth quality criterion: the surrogate is acceptable if and only if

$$\text{Var}_{q_\theta}(\hat{\phi}) \leq \frac{C}{MF_Q^*}, \quad C \gtrsim 1. \quad (11)$$

A surrogate whose posterior variance consistently exceeds $1/(MF_Q^*)$ is suboptimal and should be retrained [Cranmer et al., 2020].

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, with demonstrated results: RL has discovered 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 noise estimate.
- **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) that reduce δ_{RL} below what the analytically designed CPMG sequence achieves [Krenn et al., 2023]. Policy gradient optimization via REINFORCE yields:

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\pi}[\nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \cdot G_t], \quad G_t = \sum_{k=t}^T \gamma^{k-t} r_k. \quad (12)$$

The *RL optimality gap* is:

$$\delta_{\text{RL}} = 1 - \frac{F_Q(\rho_T)}{F_Q^*}. \quad (13)$$

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:

$$\hat{\mu}(r) = \arg \max_{\mu \geq 0} \sum_i [\mu(r_i) - n_i \log \mu(r_i)] + \lambda \|\mu\|_{\text{TV}}, \quad (14)$$

where n_i are observed photon counts and $\|\cdot\|_{\text{TV}}$ is the total variation seminorm.

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

$$\hat{c} = \arg \min_{c \geq 0} \|Kc - C\|_2^2 + \lambda \|\nabla c\|_1. \quad (15)$$

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]:

$$\hat{S} = \arg \min_{S \geq 0} \max_{u \in \mathcal{U}} \|K(u)S - C\|_2^2 + \lambda \|S\|_1, \quad (16)$$

where \mathcal{U} is an uncertainty ellipsoid around the measured wind field u .

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 systems are in clinical evaluation [Global Quantum Intelligence, 2025].

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 CO₂ over a 20-year horizon. QLM Technology received EPA Alternative Test Method qualification in 2025 [QLM Technology, 2025]. Section 6 provides a complete end-to-end treatment of this application, including an explicit classical-versus-quantum comparison via the Fisher information advantage factor in Eq. (18).

5.4 Emerging Directions

Quantum inertial navigation [Lab/prototype] based on atom interferometers offers substantially lower drift rates than classical inertial measurement units. Government-funded field trials are underway in the US and UK [PatSnap, 2026]. 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 counts individual photons backscattered from a methane plume. 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 the Beer–Lambert law:

$$\ln \frac{N_{\text{off}}}{N_{\text{on}}} = 2\sigma_{\text{CH}_4}(\lambda_{\text{on}}) \int_0^R c(r) dr, \quad (17)$$

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 SDP Layer: Measurement Design and Quantum Advantage

SDP 2 (Section 2) applies directly: the optimal POVM for estimating column concentration is photon-number-resolving detection.

Classical-versus-quantum comparison. In the low-photon-number regime $\bar{n} \ll 1$ relevant to eye-safe long-range operation, the QFI advantage of a single-photon Fock state over a coherent-state probe (classical DIAL) with the same mean photon number \bar{n} is:

$$\frac{F_Q^{\text{spf}}}{F_Q^{\text{coh}}} = \frac{1}{1 - e^{-\bar{n}}} = \frac{1}{\bar{n}} + \frac{1}{2} + \frac{\bar{n}}{12} + O(\bar{n}^2). \quad (18)$$

At $\bar{n} = 0.1$ (representative of eye-safe standoff operation) this advantage factor is approximately **10.5**: the single-photon sensor extracts roughly ten times more information per pulse than classical coherent DIAL under ideal number-resolving detection. This factor grows as pulse energy is reduced, making the quantum advantage largest precisely where eye-safe operation is most constrained.

Deployed-hardware caveat. Deployed SPADs operate in threshold (“click”) mode; Eq. (18) therefore gives an upper bound on what current hardware achieves. SDP 2 quantifies the residual *measurement loss* precisely, identifying the gap between the deployed click detector and the ideal number-resolving POVM as an actionable engineering target.

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 (14). The QCRB from SDP 2 provides a variance floor: no estimator can achieve $\text{Var}(\hat{c}) < 1/(MF_Q^*)$ for a given photon budget M . 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 (15). 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 K is stochastic:

$$\text{choose trajectory } \arg \min_{\text{traj}} \mathbb{E}_u[\text{cond}(K(u))]. \quad (19)$$

6.5 Stage 3: Source Attribution via Advection-Diffusion Inversion

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

$$\frac{\partial c}{\partial t} + \mathbf{u} \cdot \nabla c = D \nabla^2 c + S(\mathbf{x}), \quad (20)$$

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 (16) [Ben-Tal et al., 2009].

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 $-\text{Var}(S) - \lambda \cdot \text{flight cost}$ [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]:

$$\min_{\text{schedule}} \mathbb{E} \left[\sum_k w_k \cdot \mathbf{1}[\text{leak } k \text{ unrepaired}] \cdot \hat{S}_k \right] \quad \text{s.t. crew and time constraints}, \quad (21)$$

where w_k weights leaks by global warming potential and \hat{S}_k is the estimated flux with its posterior distribution propagated from Stage 3. This is a stochastic integer program consuming the pipeline output [Shapiro et al., 2009].

6.8 Summary: The Full Stack

Table 4 summarizes the complete methane monitoring pipeline.

7 Benchmarking and Evaluation

The SDP framework provides a principled upper bound on achievable precision—separating hardware limitations from algorithmic inefficiency, diagnosing precisely where performance is being lost (in

Table 4: The methane LiDAR end-to-end stack. Each row is one layer of the optimization–inference pipeline; the SDP column shows how the quantum bound enters each stage.

Layer	Physics/Engineering	OR/ML Method	SDP Role
Probe design	Photon-number state $ n\rangle$; SPAD detector	SDP 2 selects optimal POVM	F_Q^* sets variance floor; advantage factor (18)
Absorption extraction	Beer–Lambert; TC-SPC histogram	Poisson MLE (14)	QCRB bounds $\text{Var}(\hat{c})$
3-D reconstruction	Ray-integration $C = Kc$	TV regularization (15)	Holevo bound via (14)
Source attribution	Advection-diffusion PDE (20)	Robust minimax (16)	Wind UQ propagated from SDP
Scan allocation	Drone flight path	GP bandit / RL MDP (6)	SDP 1 bound caps per-step IG via (8)
Dispatch	Crew scheduling	Stochastic integer program (21)	Flux UQ from Stage 3

probe state preparation, in measurement choice, or irreversibly in decoherence), and guiding investment decisions. The gap between the achieved F_C and the optimal $F_Q^*/4$ decomposes into three components:

1. **State-preparation loss** $F_Q^* - F_Q$: how far the prepared state is from optimal. Closed by SDP 1.
2. **Measurement loss** $F_Q - F_C$: how far the chosen measurement is from optimal. Closed by SDP 2.
3. **Decoherence loss** $F_Q^0 - F_Q$: precision irreversibly lost to the environment. Not recoverable by protocol design.

These three losses connect to the RL optimality gap (13) and the surrogate quality criterion (11) by the inequality chain:

$$\underbrace{\frac{1}{MF_Q^*}}_{\text{SDP 1 bound}} \leq \underbrace{\frac{1}{MF_Q}}_{\text{actual state}} \leq \underbrace{\frac{1}{MF_C}}_{\text{actual measurement}} \leq \text{Var}(\hat{\phi}). \quad (22)$$

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 :

$$F_Q^0 = (\gamma_e \tau)^2, \quad F_Q = (\gamma_e \tau)^2 e^{-2\Gamma \tau}.$$

SDP 1 confirms that the Hadamard-prepared Ramsey state $(|0\rangle + |1\rangle)/\sqrt{2}$ is optimal under pure dephasing (state-preparation loss zero). SDP 2 confirms that σ_z readout is already optimal (measurement loss zero). The net precision loss $1 - e^{-2\Gamma \tau}$ is entirely decoherence loss—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 optimally by the two SDPs of 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**—SDP 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 (21); 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 SDPs

The two SDPs in Section 2 are the right formulations, but 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. Precisely the same structural properties appear in the sensing SDPs.

Exploiting Lyapunov–Schur structure. The block-matrix constraint in SDP 1 is a Lyapunov equation linearized via Schur complement. Tailored first-order methods can reduce per-iteration cost from $O(d^6)$ (generic SDP) toward $O(d^3)$.

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, SDP 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, SDP 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 constraint $\sum_t \tau(\xi_t) \leq T_{\max}$.

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 SDP bound on F_Q^* provides a principled metric for posterior quality (11) specific to the quantum sensing setting.

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; ideal number-resolving architectures can achieve a Fisher information advantage of approximately $10.5\times$ over coherent-state DIAL under representative eye-safe pulse energies—a quantitative benchmark traced end-to-end through the full stack of this paper.

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 SDP formulations of Section 2 serve as enabling oracle tools—given a noise model, SDP 1 returns the probe state to prepare and SDP 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 of the probe-state SDP, 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 is a bandit problem with a QFI upper bound (8); 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 5 collects the principal symbols used throughout the paper. Note: $\text{Tr}(\cdot)$ denotes the matrix trace (sum of diagonal entries) throughout.

Table 5: 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}	Decoherence map (CPTP)
$\{E_i\}$	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 shots
$\widehat{\Delta}\phi_{\text{SQL}}$	Standard quantum limit
$\widehat{\Delta}\phi_{\text{HL}}$	Heisenberg limit
Σ	Estimator covariance matrix
W	Parameter weight matrix
δ_{RL}	RL optimality gap, Eq. (13)
\bar{n}	Mean photon number per pulse

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

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 $\theta = t(\Delta E + \gamma V_0)/\hbar$.

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

$$p_{0 \rightarrow 1} = \frac{1}{2} \left[1 - \cos \frac{t(\Delta E + \gamma V_0)}{\hbar} \right]. \quad (23)$$

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

$$F_C = \frac{(\partial_{V_0} p_{0 \rightarrow 1})^2}{p_{0 \rightarrow 1}(1 - p_{0 \rightarrow 1})} = \left(\frac{t\gamma}{\hbar} \right)^2. \quad (24)$$

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, the central limit theorem gives the *standard quantum limit* (SQL):

$$\widehat{\Delta\phi}_{\text{SQL}} = \frac{1}{\sqrt{NM}}. \quad (25)$$

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):

$$\widehat{\Delta\phi}_{\text{HL}} = \frac{1}{N\sqrt{M}}, \quad (26)$$

a quadratic improvement in 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, $F_C = (\gamma_e\tau/\hbar)^2 = F_Q$, confirming optimality. The three-component gap decomposition of Section 7 follows from the inequality chain (22):

$$\underbrace{\frac{1}{MF_Q^*}}_{\text{SDP 1 bound}} \leq \underbrace{\frac{1}{MF_Q}}_{\text{actual state}} \leq \underbrace{\frac{1}{MF_C}}_{\text{actual meas.}} \leq \text{Var}(\hat{\phi}).$$

For the NV center under pure dephasing, SDP 1 and SDP 2 both confirm 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. Government investment has been substantial: 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 have collectively committed significant funding [DataIntelco, 2026]. A meaningful fraction is now directed toward field trials rather than laboratory demonstrations, reflecting increasing hardware maturity. Commercial firms include AOSense, Exail (formerly Muquans), and GEM Systems (gravimetry, atomic clocks, magnetometry); QLM Technology and Bridger Photonics (single-photon LiDAR gas detection); and Q-CTRL (quantum control software). Market projections should be interpreted cautiously, since definitions vary across analyst reports [Global Market Insights, 2025].

D Glossary of Acronyms

Table 6: 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
NPA	Navascués–Pironio–Acín (SDP hierarchy)
NV	Nitrogen-Vacancy (center in diamond)
OPM	Optically Pumped Magnetometer
POVM	Positive Operator-Valued Measure
PPT	Positive Partial Transpose (entanglement criterion)
Q-INS	Quantum Inertial Navigation System
QCRB	Quantum Cramér–Rao Bound
QEC	Quantum Error Correction
QFI	Quantum Fisher Information
QFT	Quantum Fourier Transform
QPE	Quantum Phase Estimation
RL	Reinforcement Learning
SBI	Simulation-Based Inference
SLD	Symmetric Logarithmic Derivative
SPAD	Single-Photon Avalanche Detector
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.
- DataIntel (2026). NV center diamond for quantum sensing market research report 2034. Technical report, DataIntel.
- 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.
- Global Market Insights (2025). Quantum sensors market size, share and industry report 2035. Technical report.
- Global Quantum Intelligence (2025). Top predictions for quantum technology in 2026. Technical report.
- Hayashi, M. and Ouyang, Y. (2023). Tight Cramér–Rao bound for multiparameter estimation via semidefinite programming. *Physical Review A*, 108:062609.
- 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.
- MIT News Office (2025). LiDAR helps gas industry find methane leaks and avoid costly losses. <https://news.mit.edu/2025/>.
- Muquans/Exail (2021). Absolute quantum gravimeter (AQG). Technical report, Exail. <https://www.exail.com/>.
- Nature Collections (2025). Harnessing quantum: Progress towards real-world applications. <https://www.nature.com/collections/fdjhbfihibi>.
- Nielsen, M. A. and Chuang, I. L. (2010). *Quantum Computation and Quantum Information*. Cambridge University Press, 10th anniversary edition.
- PatSnap (2026). Quantum sensing technology landscape 2026. Technical report.

- 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.
- Shapiro, A., Dentcheva, D., and Ruszczyński, A. (2009). *Lectures on Stochastic Programming: Modeling and Theory*. SIAM/MPS.
- 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.