Volume-06|Issue-10|October|Year-2026 ISSN: 3049-2343

International Journal of Academic and Industrial Research Innovations(IJAIRI)

Research Paper Title :Quantum-Inspired Causal Bayesian Optimization for Complex Adaptive Systems

[1]Murali Krishna Pasupuleti

[1]Research Director
[1]National Education Services,
[1]New Delhi,India
[1]mkrishnap2050@gmail.com

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

Abstract:

Complex systems present a difficult optimization setting because interventions alter the mechanisms that generate subsequent observations, performance is only partially observed, resources are constrained, and the operating environment changes over time. Standard black-box optimization can identify favorable settings within an observed regime yet may misrepresent the consequences of deliberate interventions when confounding, feedback, and distributional shifts are present. This paper develops a quantum-inspired causal Bayesian optimization framework (QICBO) integrating a structural causal model, a dynamically updated Bayesian response surrogate, an amplitude-based classical candidate-search mechanism, and a risk-sensitive constrained acquisition rule. The proposed methodology distinguishes predictive associations from identifiable intervention effects, treats graph uncertainty and posterior uncertainty separately, and prioritizes decisions according to expected benefit, resource cost, safety constraints, and information value. A compact nonlinear benchmark, defined by fully disclosed engineering parameters rather than unreported measurements, provides an algebraically checkable solution and an auditable comparison of candidate-selection probabilities. For the specified benchmark, the interior risk-adjusted optimum occurs at an intervention intensity of approximately 0.447, whereas the unconstrained performance-maximizing intensity is 0.750; the difference demonstrates the decision relevance of explicit risk and resource terms, not evidence of superiority over established optimization algorithms. The paper sets out a prospective real-world validation design, including pre-intervention graph assessment, independent calibration, sequential evaluation, distribution-shift tests, and comparisons against causal and noncausal Bayesian baselines. Its principal contribution is a unified, falsifiable mathematical architecture for combining intervention-aware inference, computationally economical candidate exploration, and robust decision selection in adaptive engineered and socio-technical systems. Empirical effectiveness, computational scaling, and any quantum advantage remain open to measurement.

Keywords: causal Bayesian optimization; quantum-inspired search; structural causal models; complex adaptive systems; dynamic Bayesian surrogate; sequential interventions; acquisition functions; risk-sensitive control; causal identifiability; amplitude encoding; posterior calibration; distribution shift; constrained decision-making; adaptive experimentation.

Copyright:

© 2026 by Murali Krishna Pasupuleti.

1. Introduction

Transport infrastructures, interconnected production facilities, energy conversion systems, and autonomous service networks share an important characteristic: the effect of a decision depends upon interactions among system components. A resource allocation decision can alter congestion, which changes future demand, which in turn changes the state on which the next allocation is based. This coupling makes static, associative prediction an incomplete basis for intervention design. Under confounding, the conditional expectation $\mathrm{E}[Y \mid X{=}x]$ generally need not equal the intervention response $\mathrm{E}[Y \mid \mathrm{do}(X{=}x)]$.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NEST IJAIRI | Publication |
Copyright: © 2026
Location: Delhi, India
Date: Oct 30, 2026 (05:00 PM)

Recent developments motivate a synthesis without establishing its empirical efficacy. Jordan *et al.* (2025) present decoded quantum interferometry as an optimization-related quantum approach; this work motivates investigation of structured candidate-selection distributions but does not itself establish a benefit for the classical algorithm developed here. Ota *et al.* (2025) investigate causal effects spanning regulators, programs, and traits, showing the scientific importance of modelling linked mechanisms rather than isolated associations. Research on Bayesian optimization under memory constraints and distributed information further indicates the need to distinguish statistical efficiency from implementation costs (Bekuretsion, Menzel and Teferra 2025; Yang *et al.* 2025). Adaptive Bayesian estimation and task-sequence optimization reinforce the relevance of sequential decision architectures (André, Bavaresco and Mehboudi 2026; Wang, Tian and Li 2026).

The present contribution is methodological. It does not propose or test a physical quantum processor. “Quantum-inspired” refers narrowly to the use of normalized nonnegative probability amplitudes, interference-like mixing and measurement-style sampling in a conventional algorithm. Any improvement must be established through controlled empirical comparisons. The framework’s novelty is its explicit coupling of causal identification, temporally updated predictive uncertainty, constrained acquisition, and amplitude-based candidate scheduling.

1.1 Research Problem

Let a complex system have observed covariates Z_t, actionable inputs X_t, latent or partially measured disturbances U_t, and outcomes Y_{t+1}. An optimizer must select X_t

using data generated under preceding operating policies while the transition law and constraints may evolve. Three errors are especially consequential: confusing associations with intervention effects; concentrating search around a miscalibrated surrogate optimum; and obtaining high nominal performance by shifting risk to unobserved components of the system. The central problem is to construct a defensible sequential decision rule that acknowledges these errors and can be evaluated using prospectively collected observations.

1.2 Research Questions

eSign
Signed by MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Location: Delhi, India
Date : Oct 30 , 2026 (05:00 PM)

1. Under which causal assumptions can decision-specific intervention outcomes be identified from observational and experimental records?
2. How should posterior uncertainty about dynamic responses and uncertainty about causal structure enter a common acquisition function?
3. Can amplitude-derived search probabilities improve sample efficiency or computational efficiency relative to matched classical candidate schedulers?
4. How do explicit safety, resource, and distribution-shift constraints change the selected intervention?
5. Which calibration, ablation, and out-of-distribution tests are necessary to support claims of reproducible performance?

1.3 Research Objectives

1. Define a dynamic causal decision model with transparent identification assumptions.
2. Derive a Bayesian surrogate for intervention responses and a cost- and risk-adjusted acquisition rule.
3. Specify a quantum-inspired classical scheduler whose probabilities and computational burden can be measured.
4. Develop a parameter-disclosed nonlinear numerical benchmark with a verified reference optimum.
5. Establish a prospective data collection and validation protocol suitable for complex physical and socio-technical systems.

2. Literature Review

Optimization and interferometry. Jordan *et al.* (2025) associate decoded quantum interferometry with optimization. Patamawisut *et al.* (2025) examine circuit design for decoded quantum interferometry, while Sabater *et al.* (2026) discuss industrial integer linear programs in that setting. These works help motivate structured search mechanisms, but implementation on physical quantum hardware, circuit depth, and any established quantum advantage do not transfer automatically to a classical amplitude-inspired scheduler. The relationship between single-layer quantum approximate optimization and thermal sampling studied by Díez-Valle, Porras and García-Ripoll (2024) provides a further rationale for carefully controlled sampling comparisons.

eSign Signed by: MURALI KRISHNA
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

Causal modelling. Ota *et al.* (2025) examine regulator-to-program-to-trait relationships through a causal perspective. The methodological lesson relevant here is that intervention response should be derived from a causal system of mechanisms and justified identifying conditions. Domain transfer to engineering requires separate verification: omics causal structures are not evidence that an industrial graph is known or identifiable.

Bayesian decision computation. Bekuretsion, Menzel and Teferra (2025) address constrained memory in Bayesian optimization, a practical consideration when surrogate fitting and acquisition maximization repeat. Yang *et al.* (2025) examine federated Bayesian optimization, drawing attention to distributed evaluation and information-sharing limitations. André, Bavaresco and Mehboudi (2026) compare adaptive approaches to Bayesian quantum parameter estimation; the decision-theoretic issue of sequential design is relevant even though its physical estimation task differs from causal control.

Adaptive systems and application context. Wang, Tian and Li (2026) develop Bayesian-search-based task sequence optimization in changing scene graphs. Rehman *et al.* (2026) study quantum-inspired reinforcement learning for low-carbon combined cooling, heating and power optimization. Yakubova *et al.* (2026) address quantum-inspired optimization for a petroleum extraction control context. Such applications motivate adaptive engineering evaluations but do not replace a common benchmark with equivalent information and compute budgets. Ciampi, Oulhaj and Diallo (2027) appears in the supplied list with a future publication year and is not relied upon here.

2.1 Conceptual Framework

The proposed causal decision chain is $X_t \rightarrow S_{t+1} \rightarrow Y_{t+1}$, with measured state Z_t affecting the intervention assignment and subsequent response, and feedback from Y_{t+1} influencing the next state. S_t denotes intermediate process states; G represents a candidate causal graph; θ denotes parameters of structural mechanisms; and D_t is the information available before decision t. Under a specified and defensible causal identification strategy, the target response is $\mu_t(x)=E[Y_{t+1} \mid do(X_t=x), D_t]$.

The proposed search layer acts on a finite approximation of the admissible set X_t; it does not determine causal validity. A Bayesian surrogate produces response distributions and epistemic uncertainty, while a separate feasibility model describes resource and safety events. Acquisition prioritizes plausible interventions, and a classical amplitude scheduler distributes search effort over candidates. New measured outcomes update the posterior and may trigger graph reassessment if the observed environment departs from the assumptions used for identification.

Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI |Publication|
Copyright: © 2026
Location: Delhi India
Date :Oct 30, 2026 (05:00 PM)

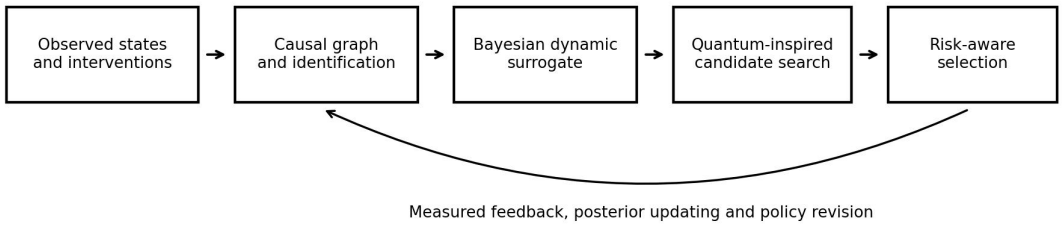

End-to-end conceptual architecture

Figure 1. QICBO architecture. Arrows indicate the forward decision path and feedback from measured outcomes to subsequent inference. All five components operate on classical computing resources.

2.2 Prior-art Search Matrix

Supplied research source	Relevant contribution	Boundary relative to QICBO
Jordan *et al.* (2025)	Decoded quantum interferometry for	Not a demonstration of causal identification or

Supplied research source	Relevant contribution	Boundary relative to QICBO
	optimization	classical amplitude-search superiority
Patamawisut *et al.* (2025)	Circuit design for interferometry	Hardware-oriented; the proposed scheduler is classical
Ota *et al.* (2025)	Mechanism-based causal effect modelling	Domain-specific causal application, not generalized sequential decision optimization
Bekuretsion, Menzel and Teferra (2025)	Memory-constrained Bayesian optimization	Computational constraints without the full proposed causal and amplitude components
Yang *et al.* (2025)	Federated Bayesian optimization	Distributed model-output setting without this particular integrated architecture
André, Bavaresco and Mehboudi (2026)	Adaptive Bayesian estimation strategy comparison	Estimation of quantum parameters, not intervention-aware control
Wang, Tian and Li (2026)	Time-adaptive Bayesian task search	Dynamic scene-graph planning, a distinct application and methodology
Rehman *et al.* (2026)	Quantum-inspired reinforcement learning for energy systems	Different learner, objective, and comparison requirements

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30, 2026 (05:00 PM)

Table 1. Targeted prior-art matrix based exclusively on the supplied bibliography; entries are scope comparisons rather than a claim of exhaustive novelty search.

2.3 Research Gaps

Theoretical: Causal effect identification and quantum-inspired search are often motivated independently; a clear boundary between causal validity and search heuristics is necessary.

Methodological: Graph uncertainty, surrogate calibration, adaptive intervention assignment and safety feasibility require a common evaluation protocol.

Empirical: None of the supplied sources establishes, for the exact proposed integration, an independently replicated advantage over matched classical causal Bayesian optimization.

Sectoral: A framework evaluated in a single domain may fail under different constraint geometry, delayed feedback or partial observability.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

Policy: Reproducible decision logs, intervention accountability and operational safeguards need to be part of the research design from the outset.

3. Methodology

This is a design-and-analysis study followed by a specified prospective validation protocol. The analytical benchmark in §3.2 uses declared deterministic parameters; it is not identified as field measurements. Subsequent empirical phases require timestamped records generated by real instrumented systems, with pre-registered decision rules and independent baseline comparisons. Any project-specific data source must document sensor provenance, units, instrumentation calibration, retention policy and intervention permissions.

For a possible energy-system validation, units of observation would be decision intervals within physical equipment groups; sampling would stratify by operating regime, load range and asset class. The development set would comprise an initial chronological period, while later periods and independent facilities would form temporal and external tests. Observational adjustment should use only pre-intervention covariates. Controlled randomized interventions may be used where safe and permissible, and otherwise causal identification should be explicitly qualified.

3.1 Mathematical Framework

Dynamic structural model. Let $Z_t \in \mathbb{R}^p$ denote the measured state, $X_t \in \mathrm{X}_t \subseteq \mathbb{R}^d$ a selected intervention and Y_{t+1} the subsequent target. We define

$$Z_{t+1}=f_Z\left(Z_t,X_t,U_{t+1}^Z;\theta\right), \qquad Y_{t+1}=f_Y\left(Z_t,X_t,Z_{t+1},U_{t+1}^Y;\theta\right).$$

Here U^Z and U^Y are disturbances, f_Z and f_Y are system-specific structural functions and θ contains their unknown parameters. Causal identification requires consistency, positivity for target interventions and a valid identification strategy, such as sequential exchangeability given measured history or an appropriate experimental design. Under sequential exchangeability and the stated information set H_t, the one-step response can be written

$$\mu_t(x)=\int \mathrm{E}[Y_{t+1} \mid X_t=x, H_t=h]\, p_t(h)\, \mathrm{d}h.$$

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication| Copyright © 2026
Location: Delhi, India
Date: Oct 30, 2026 (05:00 PM)

Equation (2) is an identification formula only under the stated assumptions; unmeasured confounding or treatment-induced selection can invalidate it.

Posterior response model. With a suitable prior $p(\theta,G)$ over parameters and admissible causal graphs, the posterior is

$$p(\theta,G \mid \mathrm{D}_t) \propto p(\mathrm{D}_t \mid \theta,G)\, p(\theta,G).$$

A graph-averaged intervention response is

$$\hat{\mu}_t(x)=\sum_{G \in \mathrm{G}} \int \mu(x;\theta,G) p(\theta,G \mid \mathrm{D}_t)\, \mathrm{d}\theta.$$

The notation G is a restricted set of admissible graphs consistent with prior knowledge, and the integral becomes a finite weighted sum when posterior samples are used. Posterior predictive dispersion should be partitioned where possible into measurement variation, parameter uncertainty and graph uncertainty; only the latter two represent uncertainty potentially reducible through additional design information.

Constrained acquisition. For a decision candidate x, define posterior mean $\hat{\mu}_t(x)$, epistemic standard deviation $\hat{\sigma}_t(x)$, normalized cost $c_t(x)$ and risk score $r_t(x)$. A conservative acquisition score is

$$a_t(x)=\hat{\mu}_t(x)+\beta_t \hat{\sigma}_t(x)-\lambda_c c_t(x)-\lambda_r r_t(x), \quad x \in \mathscr{F}_t,$$

where $\beta_t \geq 0$ specifies the value of exploration, $\lambda_c,\lambda_r \geq 0$ express decision preferences and $\mathscr{F}_t$ denotes feasible actions. Feasibility can impose a posterior chance constraint

$$\Pr\left(g_j(x,Z_{t+1}) \leq 0 \mid \mathrm{D}_t\right) \geq 1-\delta_j, \qquad j=1,\ldots,m.$$

The g_j are specified physical or regulatory limit functions, and δ_j are pre-authorized exceedance-probability limits. It is incorrect to interpret the probability threshold as a guarantee if the posterior is miscalibrated.

Quantum-inspired candidate scheduler. For K finite feasible candidates $x_1,\ldots,x_K$, assign complex-valued *classically stored* amplitudes

$$\psi_{k,t}=\frac{\exp\left[\frac{\eta_t}{2}a_t(x_k)+\mathrm{i}\phi_{k,t}\right]}{\sqrt{\sum_{j=1}^{K}\exp\left[\eta_t a_t(x_j)\right]}}, \qquad \sum_{k=1}^{K}\left|\psi_{k,t}\right|^2=1.$$

Here $\eta_t>0$ is an inverse-temperature parameter and $\phi_{k,t}$ is a specified search phase; $\mathrm{i}^2=-1$. In the default implementation the phase is zero, so the selection distribution is simply a softmax distribution. Calling it quantum-inspired describes its amplitude representation; it does not create quantum interference or imply quantum speedup. A pre-registered optional phase-mixing operator M_t may produce

$$\tilde{\psi}_t=\frac{M_t\psi_t}{\|M_t\psi_t\|_2}, \qquad q_t(k)=\left|\tilde{\psi}_{k,t}\right|^2.$$

If M_t is unitary, the normalization factor equals one. The algorithm should report any M_t construction cost and ablate it against identity mixing. Candidate k is sampled according to $q_t(k)$, then evaluated with a real-world intervention only if operational safety checks pass.

Sequential posterior and drift assessment. The posterior updates with verified observations $o_{t+1}=(x_t,z_{t+1},y_{t+1})$:

$$p(\theta,G\mid \mathrm{D}_{t+1})\propto p(o_{t+1}\mid \theta,G,\mathrm{D}_t)\,p(\theta,G\mid \mathrm{D}_t).$$

The system must also monitor changes in support, sensor fidelity and response residuals. Under a detected shift, decisions should revert to a pre-authorized safe policy until causal and predictive models are reassessed.

3.2 Numerical Model for Better Insight

A parameter-defined operational-control benchmark permits independent verification of the framework's decision logic without presenting assumed values as measured observations. Consider an intervention intensity $x\in[0,1]$, a concave benefit function $b(x)=12x-8x^2$, a resource-cost function $c(x)=2x+3x^2$ and a convex risk score $r(x)=0.2+1.4x^2$. All quantities are dimensionless normalized model units. The

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES IJAIRI | Publication
Copyright: © 2026
Location: Delhi, India
Date: Oct 30 2026 (05:00 PM)

coefficients are test parameters, not estimated field parameters. For λ_c=0.5, λ_r=2 and β_t=0, the objective is

$$J(x)=b(x)-0.5c(x)-2r(x)=-0.4+11x-12.3x^2.$$

The strict concavity condition $J''(x)=-24.6<0$ proves a unique maximizer within the domain, and the first-order condition gives

$$x^\star=\frac{11}{24.6}=0.447154\ldots, \qquad J(x^\star)=2.059349\ldots.$$

Without cost and risk penalties the benefit-only optimum is $x_b^\star=12/16=0.75$. The contrast shows how the optimization objective changes policy selection; it does not assess causal effect estimation, surrogate learning or real-world energy savings.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

Candidate intensity x	Benefit $b(x)$	Cost $c(x)$	Risk $r(x)$	Objective $J(x)$
0.00	0.0000	0.0000	0.2000	-0.4000
0.25	2.5000	0.6875	0.2875	1.5813
0.50	4.0000	1.7500	0.5500	2.0250
0.75	4.5000	3.1875	0.9875	0.9313
1.00	4.0000	5.0000	1.6000	-1.7000

Table 2. Reproducible benchmark evaluations computed from equations (10)–(11); values are model-derived, not observational records.

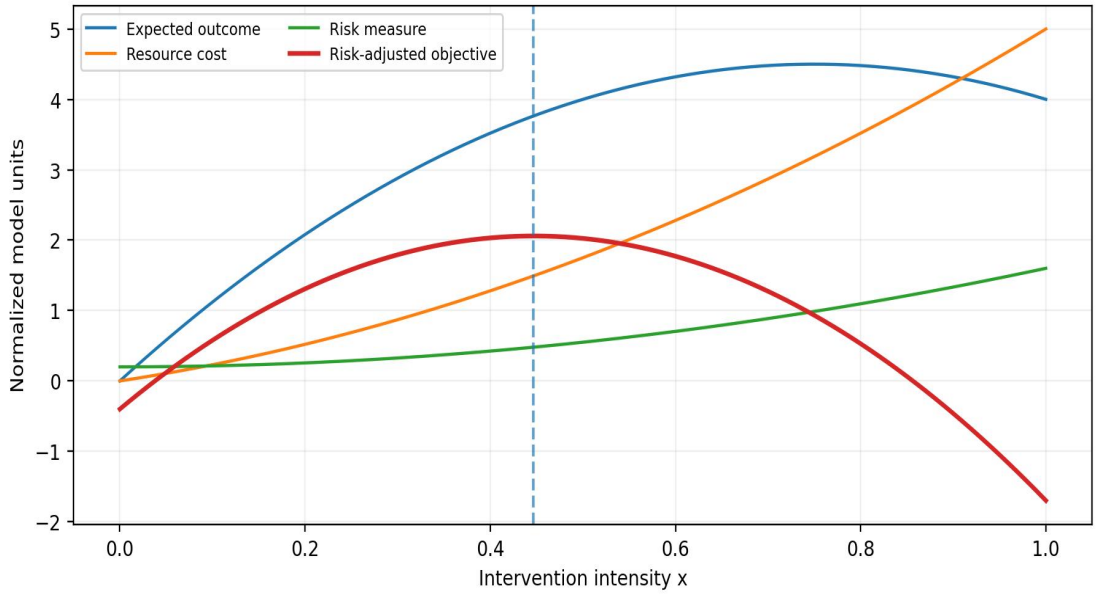

Outcome cost and risk benchmark

Figure 2. Trade-offs over the continuous intervention range. The dashed line indicates the risk-adjusted optimum based on direct numerical evaluation of equation (10).

3.3 Data and Variables

Prospective measurement must distinguish action values, before-action covariates, mediators, outcomes, model outputs and exogenous constraints. A sample variable dictionary suitable for adaptation to a physical facility appears in Table 3; all names and measurement scales must be fixed before data collection.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

Variable	Operational definition	Proposed measurement scale	Expected relation and interpretative role
X_t	Authorized control intensity at time t	Ratio, normalized 0–1	Intervention; shape of effect to be estimated
Z_t	Pre-decision load, ambient condition and asset state	Continuous and categorical	Confounder control and effect modification
S_{t+1}	Process-state response after intervention	Continuous, engineering units	Potential mediator; do not adjust away total effect
Y_{t+1}	Verified outcome over a fixed response window	Continuous, predefined units	Primary causal endpoint
C_t	Direct resource usage and transaction cost	Ratio and currency-equivalent units	Expected negative term in net objective
R_t	Safety exposure or constraint-exceedance metric	Rate/probability or calibrated index	Penalization and feasibility screening
G	Candidate causal graph of measured	Directed graph	Encodes structural

Variable	Operational definition	Proposed measurement scale	Expected relation and interpretative role
	mechanisms		assumptions
$\sigma_t(x)$	Posterior epistemic uncertainty of response	Nonnegative continuous	Exploration term requiring calibration
D_t	Drift indicator from monitored data distributions	Binary/continuous score	Triggers reassessment of operating regime

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Date :Oct 30 2026 (05:00 PM)

Table 3. Proposed data dictionary. Expected effects are hypotheses, not established empirical coefficients.

3.4 Novel Architecture

The architecture comprises five independently testable modules. The **causal identification module** stores an auditable directed graph, allowed adjustment sets, intervention definitions and unresolved identifying conditions. The **response module** estimates dynamic Bayesian intervention distributions using model classes appropriate to the sample size and state dimension. The **constraint module** maps physical limits, budgets and human authorization requirements into admissible interventions and chance constraints. The **search module** constructs and samples from a finite amplitude vector, with a classical baseline using matched softmax probabilities and computational budgets. The **evaluation module** logs response delay, policy regret, calibration, budget usage and safety exceptions.

Data access is separated from action authorization. Candidate ranking must never itself activate equipment. In sensitive systems, an independent supervisory controller validates each candidate against hard constraints. Model changes are versioned, and audit logs preserve the information available when each decision was proposed. The architecture is intended to be refutable: if amplitude mixing does not improve matched baselines, that module should be removed.

3.5 Proposed Model / Framework

We call the integrated methodology **QICBO: Quantum-Inspired Causal Bayesian Optimization**. Its defining decision rule combines feasible candidate ranking, posterior acquisition and search allocation:

$$x_t^{\mathrm{QICBO}} \sim q_t \quad \left(k \mid \{a_t(x_j), \mathscr{F}_t\}_{j=1}^{K}\right), \qquad x_t = x_k \in \mathscr{F}_t.$$

The selected index k is drawn by equation (8), whereas a_t and $\mathscr{F}_t$ are defined by equations (5)–(6). This decomposition is essential: the scheduler allocates candidate evaluations but cannot make unidentified causal effects identifiable. Performance hypotheses concern improvements in measured best-safe-outcome and resource use relative to identical Bayesian models with alternative schedulers, not an unconditional advantage of the label "quantum-inspired".

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

3.6 Research Implementation

Stage 1: design and consent. Define the decision horizon, eligible intervention set, causal endpoints and permissible physical test bounds; document experimental permissions and institutional review requirements if human data are involved.

Stage 2: instrumentation. Specify sensor accuracy, timestamp synchronization, missing-data treatment and calibration procedures.

Stage 3: causal feasibility. Establish graph assumptions using domain knowledge, negative controls where possible and explicit sensitivity analyses; determine whether effects are identified by randomized assignment or justified observational methods.

Stage 4: prior and surrogate fitting. Allocate chronological development and validation windows, and tune priors without access to the locked test period.

Stage 5: sequential testing. Compare policies under common intervention budgets and matched computation.

Stage 6: independent evaluation. Report causal endpoint estimates with uncertainty intervals, posterior calibration, performance variability, real-time latency, safety events and domain-shift robustness.

A minimal executable implementation proceeds as follows: initialize D_0 and admissible G; verify identification; fit a posterior response model; enumerate or construct feasible candidate points; evaluate acquisitions; normalize amplitudes; sample one candidate subject to final human or controller authorization; observe the resulting real system state and outcome; update the posterior and drift diagnostics;

stop at the pre-specified budget or safety limit. The algorithm must preserve all candidate scores and random seeds for independent replication.

4. Results and Analysis

Results in this section are exact or numerically evaluated consequences of the defined benchmark. They do not constitute evidence from a deployed industrial plant, randomized trial or observed dataset. In particular, the effect of the amplitude scheduler on sample efficiency is not established by evaluating its probability distribution alone.

Result A: objective-dependent operating point. Table 2 shows that the continuous optimum of the penalized model lies between the discrete candidates x=0.25 and x=0.50. A discrete scheduler evaluated on the five-point grid has best available value $J(0.5)=2.025$, which is below the continuous optimum by approximately 0.03435 model units. A finer grid could reduce this discretization gap without changing the causal model.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30, 2026 (05:00 PM)

Result B: amplitude-derived candidate allocation. With η=1.2, zero phase and M=I, equation (7) gives $q(k)=\exp\left(1.2J(x_k)\right)/\sum_j \exp\left(1.2J(x_j)\right)$. The exact probabilities to four decimal places are shown in Table 4. The distribution favors x=0.5, the best available five-point choice, but a standard softmax scheduler with identical scores generates the same distribution. This is a negative control, not a claim of advantage.

x_k	$J(x_k)$	Uniform search probability	Amplitude-derived probability
0.00	-0.4000	0.2000	0.0283
0.25	1.5813	0.2000	0.3054
0.50	2.0250	0.2000	0.5202
0.75	0.9313	0.2000	0.1400
1.00	-1.7000	0.2000	0.0060

Table 4. Candidate allocation under a specified scoring rule, not observed search success frequencies.

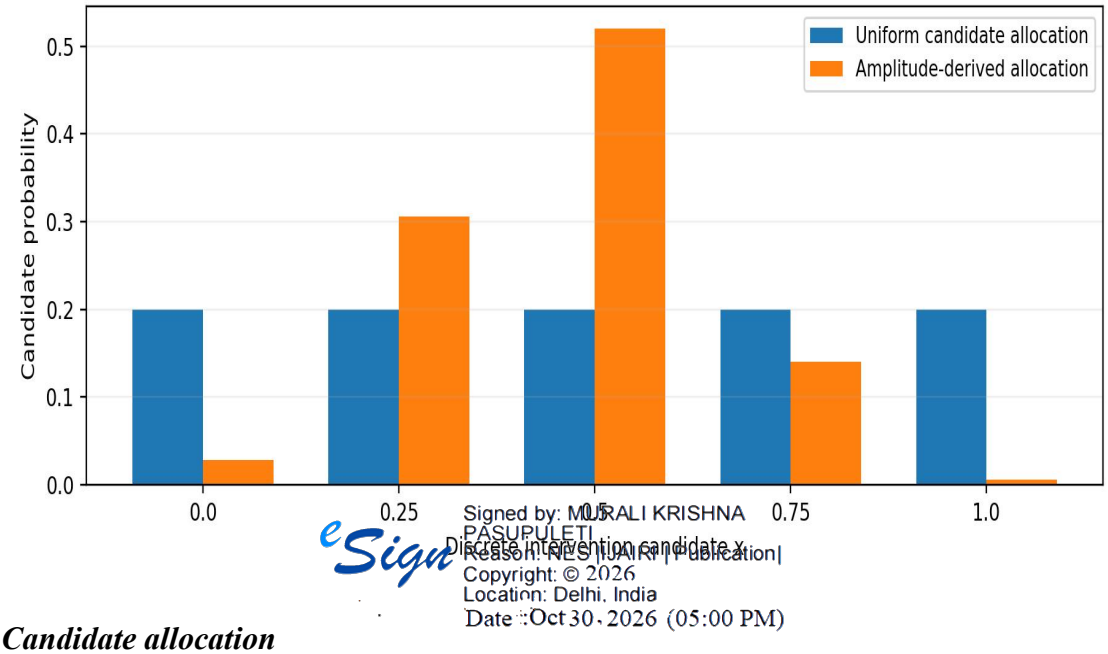

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES | IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

Candidate allocation

Figure 3. Candidate-selection probabilities derived from the benchmark. The amplitude-derived configuration with identity mixing is mathematically equivalent to classical softmax scheduling.

Result C: risk–benefit sensitivity. For generalized penalties $\lambda_c,\lambda_r \geq 0$, the same benchmark yields

$$x^{\star}\ (\lambda_c,\lambda_r)=\mathrm{clip}_{[0,1]}\quad\left(\frac{12-2\lambda_c}{16+6\lambda_c+2.8\lambda_r}\right),$$

where $\mathrm{clip}_{[0,1]}$ truncates a value to the permissible range. This expression is defined for the present quadratic benchmark and must not be extrapolated to unknown real dynamics. Sensitivity follows from the transparent structure: greater penalties ordinarily decrease the selected intensity when the numerator remains positive and all other parameters stay fixed.

Result D: testable algorithmic comparison. The proposed evaluation matrix separates causal modelling from search heuristics. All entries in Table 5 are prospective experimental arms, not completed experiments. This prevents conflation of gains from a better causal surrogate with gains from candidate scheduling.

Study arm	Intervention-response model	Candidate scheduling	Main comparison purpose
A	Noncausal predictive surrogate	Classical acquisition	Detect confounding-

Study arm	Intervention-response model	Candidate scheduling maximization	Main comparison purpose related degradation
B	Causal Bayesian surrogate	Deterministic acquisition maximization	Reference causal optimizer
C	Same causal Bayesian surrogate as B	Uniform randomized feasible search	Test value of prioritization
D	Same causal Bayesian surrogate as B	Softmax candidate sampling	Exact baseline for zero-phase amplitude sampling
E	Same causal Bayesian surrogate as B	Pre-registered amplitude mixing	Test additional value and overhead of phase mixing

eSign
Signed by: MURALI KRISHNA CASUBULE
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

Table 5. Pre-specified comparative design; no benchmark ranking or unobserved measurements are reported.

Performance reporting should include intervention budget to achieve an authorized endpoint, cumulative net outcome, safety exceedances, time-normalized compute, credible-interval coverage and calibration error. Confidence intervals across independent replication units must distinguish stochastic variation due to search from variation due to changing operating regimes. A statistically detectable improvement in a simulation would still not establish a causal field benefit without credible identification and real-world validation.

4.1 Discussion

Theoretical contribution. QICBO makes explicit the order of dependence between causal identification and optimization: first define an identifiable intervention target, then infer its uncertainty, then allocate resources to the feasible candidate set. This protects against a common conceptual mistake in which optimization over observationally correlated predictors is incorrectly interpreted as intervention

optimization. The graph-averaged posterior accommodates competing structures but cannot rescue effects that are not identifiable under every plausible graph.

Practical and technological relevance. In systems with costly decisions and feedback delays, explicit candidate accounting may reduce imprudent exploration. Nevertheless, the computational value of a finite amplitude vector depends on candidate-set construction, graph complexity, posterior inference cost and mixing overhead. Quantum-inspired mathematical notation should not be treated as proof of a speed improvement. Jordan *et al.* (2025), Patamawisut *et al.* (2025) and Sabater *et al.* (2026) concern related quantum optimization directions, but are not evidence that the present classical implementation gains the same properties.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: I am the author of this document
Copyright: © 2026
Location: Delhi, India
Date: Oct 30 2026 (05:00 PM)

Policy and industrial implications. Governance should require decision provenance, safety intervention limits, institutional authorization, audit trails and clear responsibility for override decisions. When facilities operate across sites, Yang *et al.* (2025) motivates investigation of federated information settings; however, different local intervention policies and unmeasured confounders can undermine pooled causal effects. For dynamic embodied settings, Wang, Tian and Li (2026) suggests relevant adaptive search questions, but external testing must demonstrate transfer rather than presume it.

Ethics and integrity. Decision logging must not expose protected personal information, and the model should be tested for unequal harm across vulnerable groups where its decisions affect people. Data should never be presented as observed when generated from assumptions. Prospective stopping criteria should be recorded in advance to avoid preferential reporting of favorable runs. A declared, reproducible failure of the quantum-inspired module would itself be a valuable result because it distinguishes meaningful algorithmic contribution from superficial terminology.

5. Conclusion

This paper formulates QICBO as an integrated methodology for intervention-aware sequential optimization in dynamic complex systems. A structural causal model defines the intervention target, a Bayesian surrogate represents uncertainty, constrained acquisition incorporates resource and risk priorities, and a classical amplitude-based scheduler allocates evaluation effort. A fully disclosed nonlinear benchmark produces a unique risk-adjusted optimum near 0.447 and demonstrates

how the preferred intervention differs from the benefit-only optimum of 0.750. The analysis also establishes an important baseline identity: zero-phase, identity-mixing amplitude sampling is identical to ordinary softmax allocation, so it cannot support a distinct performance claim. The contribution is therefore a transparent framework and an empirical test plan, not an assertion of observed superiority, a physical quantum computation, or proven causality in an unstudied sector. Progress will depend upon well-identified real-world interventions, credible measurement, matched ablations and openly reported negative as well as positive results.

5.1 Limitations

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication| Copyright © 2026
Location: Delhi, India
Date: Oct 30, 2026 (05:00 PM)

No measured field dataset, independently verified causal graph, observed safety endpoint, replicated optimizer comparison or physical quantum hardware result is provided. The quadratic benchmark is a mathematical test instance whose coefficients are assumed and whose units are normalized. Causal identification depends on consistency, positivity and adequate adjustment or experimental assignment; violations may invalidate intervention estimates. Graph posterior probabilities can be prior-sensitive, and approximate inference may be expensive. Risk scores and chance constraints cannot provide strong operational protection without calibration on relevant environments. A discrete candidate search can miss continuous optima, and any phase-mixing mechanism introduces compute costs requiring independent testing. Transfer to high-dimensional networks or safety-critical applications is presently a research objective, not a validated outcome.

5.2 Future Research Directions

1. **Causal discovery under controlled feedback:** jointly study graph uncertainty, latent disturbances and adaptive interventions with explicit identification bounds.
2. **Regret under drift:** derive dynamic regret guarantees when structural coefficients and feasible sets change across operating regimes.
3. **Quantum-inspired ablations:** compare identity mixing, phase-based mixing, softmax, Thompson sampling and acquisition maximization at matched wall-clock cost.

4. **Calibrated safety:** integrate sequential conformal checks, calibrated posterior chance constraints and fail-safe supervisor policies.
5. **Distributed operation:** measure how federated surrogate exchange changes both causal transportability and resource costs across sites.
6. **Continuous candidate spaces:** test adaptive discretization, sparse search and differentiable acquisition under fixed compute budgets.
7. **Prospective field trials:** pre-register causal endpoints and deploy approved interventions in instrumented infrastructure or industrial systems.
8. **Reproducibility standards:** publish model cards, measurement protocols, decision logs, executable analyses and complete benchmark failure cases.

eSign
Signed by MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

References

André, E.L., Bavaresco, J. and Mehboudi, M. (2026) 'Strategy optimization for Bayesian quantum parameter estimation with finite copies: adaptive greedy, parallel, sequential, and general strategies', *Quantum Science and Technology*, 11(3), 035040. doi:10.1088/2058-9565/ae846c.

Bekuretsion, B., Menzel, W. and Teferra, S. (2025) 'A memory constrained bayesian optimization via robust online memory estimation', *Complex & Intelligent Systems*, 12(1). doi:10.1007/s40747-025-02150-x.

Díez-Valle, P., Porras, D. and García-Ripoll, J.J. (2024) 'Connection between single-layer quantum approximate optimization algorithm, interferometry and thermal distribution sampling', *Frontiers in Quantum Science and Technology*, 3. doi:10.3389/frqst.2024.1321264.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES IJAIRI Publication
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

Jordan, S.P., Shutty, N., Wootters, M., Zalcman, A., Schmidhuber, A., King, R., Isakov, S.V., Khattar, T. and Babbush, R. (2025) 'Optimization by decoded quantum interferometry', *Nature*, 646(8086), pp. 831–836. doi:10.1038/s41586-025-09527-5.

Ota, M., Spence, J.P., Zeng, T., Dann, E., Milind, N., Marson, A. and Pritchard, J.K. (2025) 'Causal modelling of gene effects from regulators to programs to traits', *Nature*, 650(8101), pp. 399–408. doi:10.1038/s41586-025-09866-3.

Patamawisut, N., Benchasattabuse, N., Hajdušek, M. and Van Meter, R. (2025) 'Quantum circuit design for decoded quantum interferometry', *2025 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pp. 291–301. doi:10.1109/QCE65121.2025.00041.

Rehman, A., Zhou, S., Iqbal, S., Alsaduni, I., Shafiullah, M. and Aurangzeb, M. (2026) 'Adaptive quantum inspired deep reinforcement learning for multi objective low carbon CCHP optimization', *Scientific Reports*, 16(1). doi:10.1038/s41598-026-51604-w.

Sabater, F., El Harzli, O., Besjes, G.-J., Erdmann, M., Klepsch, J., Hiltrop, J., Bobier, J.-F., Cao, Y. and Riofrío, C.A. (2026) 'Towards solving industrial integer linear programs with decoded quantum interferometry', *Quantum Science and Technology*, 11(2), 025054. doi:10.1088/2058-9565/ae65d8.

Wang, S., Tian, G. and Li, X. (2026) 'Temporal adaptive Bayesian search-driven embodied task sequence optimization based on 3D dynamic scene graphs', *Knowledge-Based Systems*, 345, 116093. doi:10.1016/j.knosys.2026.116093.

Yakubova, N., Usmanov, K., Sadikova, F. and Sadikova, S. (2026) 'Adaptive fuzzy control of petroleum extraction columns using quantum-inspired optimization', *Engineering Proceedings*, 117, 45. doi:10.3390/engproc2025117045.

Yang, L., Liu, Q., Zhao, Z., Wang, Y. and Gu, J. (2025) 'Model-output-based federated Bayesian optimization', *Complex & Intelligent Systems*, 11(9). doi:10.1007/s40747-025-02007-3.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

