Refining Causal Graphs Beyond Conditional Independence: Kernel Two-Sample Testing on Intervention Data

Anonymous authors
Paper under double-blind review

Abstract

The accuracy of causal effect estimation is governed by the accuracy of the underlying causal graph. Kernel-based conditional independence testing (KCIT), a standard approach to causal discovery, suffers from structural limitations that persist regardless of sample size. Edges obscured by latent confounding or collider bias go undetected, and under faithfulness violations, where direct and indirect effects cancel, edge detection becomes impossible in principle. We propose an integrated algorithm that overcomes these limitations by using an information source fundamentally different from conditional independence, the distributional shift in intervention data, measured via Maximum Mean Discrepancy (MMD). An MMD intervention effect map detects edges that KCIT misses and determines causal direction. Conditional MMD isolates direct effects by controlling for mediators, recovering edges under complete faithfulness violations across all seeds. Within the KCIT stage, we introduce quantum-inspired kernel joint encoding to improve sensitivity to multiplicative dependencies prevalent in physical systems. We formalize the conditions under which a physics simulator serves as a valid surrogate for Pearl's do-operator ($\mathcal{S}$-valid do-operator) and verify that three widely-used scientific simulators satisfy these conditions. On the same synthetic benchmark, the proposed method achieves an F1 of 0.673, outperforming IGSP (0.430) and UT-IGSP (0.400) by a wide margin. On three real-world scientific datasets (drug toxicity, ultrafast laser physics, and battery chemistry), our method correctly estimates positive causal effects for all outcome variables in settings where methods lacking directional information produce sign-reversed estimates due to overadjustment.

1 Introduction

Causal discovery is the problem of learning a directed acyclic graph (DAG) from observational and interventional data, and is a prerequisite for reliable causal effect estimation. Kernel-based conditional independence testing (KCIT (Zhang et al., 2012)), used within the PC algorithm (Spirtes et al., 2000), is a standard nonparametric approach. However, KCIT has structural limitations that persist regardless of sample size: edges obscured by latent confounding go undetected, collider conditioning introduces spurious dependencies, and when direct and indirect effects cancel (faithfulness violations), edge detection fails in principle. These limitations are inherent to the conditional independence paradigm itself, not to any particular kernel or sample size.

Overcoming these limitations requires an information source that is fundamentally different from conditional independence. Intervention data provides exactly this: by comparing the observational distribution $P(X_j|X_i = k)$ with the interventional distribution $P(X_j|\mathrm{do}(X_i = k))$, one can detect causal relationships that conditional independence testing misses. Maximum Mean Discrepancy (MMD) (Gretton et al., 2012) offers a nonparametric measure of such distributional shifts. While kernel mean embeddings have been applied to causal effect estimation such as counterfactual mean embeddings (Muandet et al., 2021) and conditional distributional treatment effects (Park et al., 2021), their systematic use for causal discovery (DAG structure learning) as a complement to conditional independence testing remains unexplored.

This paper fills this gap with an integrated algorithm that progressively refines a KCIT skeleton using MMD on intervention data. The MMD intervention effect map (Phase 2) detects edges that KCIT misses and determines causal direction. Conditional MMD (Phase 3) isolates direct effects even under complete path cancellation. Spurious edge removal (Phase 4) eliminates false edges. In scientific applications where real experiments are costly, computational simulators can generate intervention data. We formalize the conditions under which simulator outputs are equivalent to Pearl's do-operator.

A separate challenge concerns the detection of multiplicative dependencies among variables (products rather than sums), which are common in physical and chemical systems.[1] The Gaussian kernel's additive distance decomposition has low sensitivity to such structure. To address this, we introduce quantum-inspired kernel joint encoding into the KCIT stage (Phase 1), embedding variable products into quantum state phases via ZZFeatureMap (Havlíček et al., 2019).

Our contributions are threefold.

- **MMD-integrated algorithm** (Phases 1–4): A causal discovery algorithm that progressively refines an initial KCIT skeleton through MMD-based direction determination (KIIM), an MMD intervention effect map (Phase 2), conditional MMD for recovery under path cancellation (Phase 3), and spurious edge removal (Phase 4).
- **Verification of $\mathcal{S}$-valid do-operators**: We state three conditions under which a computational simulator serves as a valid surrogate for Pearl's do-operator (Pearl, 2009) and verify that three widely-used scientific simulators satisfy them, justifying the intervention data used in Phases 2–4.
- **Joint encoding**: Quantum-inspired kernel joint encoding for the KCIT stage (Phase 1), improving sensitivity to multiplicative dependencies.

The unifying design principle is progressive refinement of the KCIT skeleton via MMD. While KCIT provides an initial skeleton with structural limitations, MMD on intervention data supplies fundamentally different information, enabling missing edge detection, direction determination, and recovery from path cancellation. The resulting accurate causal graph directly governs downstream causal effect estimation accuracy, as verified on three real-world scientific datasets.

2 Related Work

We organize related work into six areas that provide context for the contributions of this paper.

2.1 Classical causal discovery: from observation to intervention

The earliest causal discovery frameworks relied solely on observational data. PC (Spirtes et al., 2000) and its extension FCI construct DAG skeletons via conditional independence testing. GES and NOTEARS (Zheng et al., 2018) realize structure learning through continuous optimization. As nonlinear extensions, additive noise models (ANMs) and post-nonlinear models (Zhang & Hyvärinen, 2009) established identifiability under the additive nonlinear forms $Y = f(X) + \epsilon$ and $Y = g(f(X) + \epsilon)$. Kalisch & Bühlmann (2007) proved the consistency of PC and developed theory for the high-dimensional, small-sample regime.

The use of intervention data resolves Markov equivalence classes that are unidentifiable from observations alone. Invariant Causal Prediction (ICP) (Peters et al., 2016) identifies causal sets from prediction invariance across multiple environments. IGSP (Wang et al., 2017) integrates observational and interventional data in a permutation-based framework, and UT-IGSP (Squires et al., 2020) extends this to unknown intervention targets. JCI (Mooij et al., 2020) provides a unified framework for observational and interventional data. More recently, Guo & Dong (2026) introduced interventional constraints as an alternative to interventional data for linear causal discovery. Bareinboim & Pearl (2016) established the theoretical foundations for data fusion and transportability. However, these intervention-driven methods use Gaussian kernels or linear

[1] The case includes the ponderomotive energy in laser physics ($U_p \propto I \cdot \lambda^2$).

models at the core of their conditional independence tests, resulting in low detection power for multiplicative dependencies among variables (verified in Benchmark 1A). Our work belongs to the lineage of intervention-driven causal discovery, but replaces the KCIT kernel with a quantum-inspired kernel, complements KCIT's limitations with an MMD-based intervention effect map (Phases 2–4), and captures multiplicative structure via joint encoding.

2.2 Nonlinear and non-additive dependencies

Two main directions exist for handling nonlinear dependencies. The first addresses additive nonlinearity. ANMs and post-nonlinear models (Zhang & Hyvärinen, 2009) establish identifiability under $Y = f(X) + \epsilon$ or $Y = g(f(X) + \epsilon)$. KCIT (Zhang et al., 2012) and HSIC (Gretton et al., 2007) detect nonlinear dependence via independence testing in an RKHS, but both rely on Gaussian kernels.

The second direction—handling multiplicative and interaction-type dependencies—remains limited. Polynomial kernels $(x \cdot x' + c)^d$ explicitly contain variable products but suffer from combinatorial explosion in higher dimensions. Factorization Machines (Rendle, 2010) approximate pairwise interactions at low rank but are predictive models with no extension to causal discovery. Tensor decomposition (Kolda & Bader, 2009) assumes multiplicative structure by design but requires prior knowledge for variable identification. In molecular activity prediction, manual construction of interaction terms is common practice. In contrast, our work uses a quantum-inspired kernel at the core of conditional independence testing for causal discovery, capturing multiplicative structure without prior knowledge (§4.5). Joint encoding (§4.5) is the key technique that enables this.

2.3 Deep learning and generative models for causal inference

Causal discovery using continuous optimization and deep learning has developed rapidly in recent years. Starting from NOTEARS (Zheng et al., 2018), DCDI (Brouillard et al., 2020) realized differentiable causal discovery from intervention data. ENCO (Lippe et al., 2022) and DECI (Geffner et al., 2022) propose causal discovery integrated with deep generative models. DiffAN (Sanchez & Tsaftaris, 2022) is a recent attempt at causal discovery using diffusion models.

In the line of work using generative models for interventional distribution estimation, Conditional Flow Matching (Lipman et al., 2023) has been reported to stabilize training through an ODE-based formulation. DEAR (Shen et al., 2022) and CASTLE (Kyono et al., 2020) propose VAE-based causal effect estimation. Double Machine Learning (DML) (Chernozhukov et al., 2018) provides a consistent estimator under confounding using observational data. These deep learning and generative model-based methods have achieved progress in both direction identification and effect estimation, but no integration with quantum-inspired kernels has been attempted. Our work replaces the KCIT kernel with a quantum-inspired kernel and uses a simulator do-operator as the source of intervention data, placing it in a complementary position to these approaches.

2.4 Quantum machine learning and quantum kernels

Quantum machine learning (QML), which applies quantum computing to classical data analysis, has been advancing steadily. Havlíček et al. (2019) proposed quantum kernel classification using ZZFeatureMap. Schuld (2021) established the RKHS foundations of quantum kernels. Liu et al. (2021) proved quantum advantage at the computational complexity level based on the discrete logarithm problem. Huang et al. (2022) demonstrated exponential quantum advantage in quantum data learning.

In causal discovery, Maeda et al. (2025) proposed qPC, replacing the Gaussian kernel in KCIT with a quantum kernel, and Kawaguchi (2023) applied quantum kernels to LiNGAM. These methods use only observational data. Our work differs by integrating intervention data via the $\mathcal{S}$-valid do-operator. The broader program of causal representation learning (Schölkopf et al., 2021) seeks to discover causal variables from low-level observations, a complementary direction to our work on causal discovery among given variables.

The tensor product structure of quantum kernels provides a structural advantage for detecting multiplicative dependencies in conditional independence testing (Maeda et al., 2025). However, exponential concentration in high dimensions (Thanasilp et al., 2024) limits their applicability. Our design concentrates joint encoding on the causal discovery stage where this advantage applies, and uses classical Flow Matching for the estimation stage.

2.5 Kernel embeddings for causal inference

Kernel mean embeddings (Muandet et al., 2017) provide a nonparametric framework for representing and comparing probability distributions in a reproducing kernel Hilbert space (RKHS). In causal inference, this framework has been applied in two directions.

The first direction is causal effect estimation. Muandet et al. (2021) proposed counterfactual mean embeddings (CME), which represent interventional distributions $P(Y|\mathrm{do}(X = x))$ as elements of an RKHS and define distributional treatment effects as the MMD between factual and counterfactual embeddings. Park et al. (2021) extended this to conditional settings, estimating heterogeneous treatment effects via kernel conditional mean embeddings. Fawkes & Shah (2024) developed doubly robust kernel statistics for testing equality of interventional outcome distributions. These methods assume a known or estimated causal graph and focus on quantifying effects along specified causal paths.

The second direction is causal discovery. Chen et al. (2014) used RKHS embeddings for bivariate causal discovery via the additive noise model. Lopez-Paz et al. (2015) proposed a data-driven approach to bivariate causal direction determination using kernel mean embeddings. Zhang et al. (2012) developed KCIT, which uses HSIC (Gretton et al., 2007) for conditional independence testing within the PC algorithm. More recently, Dhanakshirur et al. (2025) defined a structural causal model distance (SCMD) using MMD between interventional distributions, but for comparing learned DAGs rather than for discovering new edges.

Zenati et al. (2025) established estimation guarantees for doubly robust estimators of interventional mean embeddings. Our work bridges these two directions. We use MMD on intervention data not for effect estimation but for DAG structure learning: the MMD intervention effect map detects edges that KCIT misses, conditional MMD recovers edges under faithfulness violations, and MMD-based direction determination orients the graph. This progressive refinement of a KCIT skeleton via MMD is, to our knowledge, the first systematic integration of kernel two-sample testing into the causal discovery pipeline.

2.6 Simulator-based inference and intervention data

Research integrating physics simulators into machine learning inference has developed rapidly in recent years. Simulation-based inference (Cranmer et al., 2020; Deistler et al., 2025) is a general framework that estimates conditional densities via neural networks for physics models whose likelihood functions are analytically intractable. Likelihood-free inference and Approximate Bayesian Computation (Cranmer et al., 2020) form a related family of methods.

In the context of causal inference, attempts to use simulators as do-operators have emerged. Methods have been proposed that learn interventional distributions generated by simulators using conditional generative models (Flow Matching (Lipman et al., 2023), Normalizing Flows, etc.) and estimate causal effects. These methods have been applied to domains such as time-series data, PDE-based simulations, and chemical reactions. Integration with Neural Operators (Li et al., 2021) to streamline causal analysis of physical simulations has also been explored. While these prior works demonstrate the use of simulators as do-operators in specific cases, they do not formally define the conditions under which a simulator functions as an $\mathcal{S}$-valid do-operator. Our verification of $\mathcal{S}$-validity (§4.2) differs from these works by specifying three conditions that guarantee equivalence between simulator intervention and Pearl's do-operator. We note that simulator-generated intervention data can also be viewed through the lens of transportability (Bareinboim & Pearl, 2016), where results from a source domain (simulator) are transferred to a target domain (real system). Our $\mathcal{S}$-validity conditions can be interpreted as sufficient conditions for such transportability when the simulator faithfully implements the causal mechanisms.

3 Problem Statement

We identify three concrete challenges that this paper addresses.

Challenge 1: Structural limitations of KCIT. Kernel-based conditional independence testing (KCIT) is widely used for skeleton discovery in the PC algorithm (Spirtes et al., 2000), yet it suffers from structural limitations that persist regardless of sample size, including spurious edges due to latent confounding, conditioning bias at colliders, and cancellation between direct and indirect effects (faithfulness violations). Because these limitations are inherent to the principle of conditional independence testing itself, complementary information from a fundamentally different source is required.

Challenge 2: Procurement and validity of intervention data. Intervention-based methods (IGSP (Wang et al., 2017), UT-IGSP (Squires et al., 2020)) use intervention data but rely on Gaussian invariance tests, resulting in low detection power for multiplicative dependencies (Recall = 0.33 on Benchmark 1A; Table 4). In the physical and chemical sciences, real intervention experiments are costly, and computational simulators offer a viable alternative source of intervention data. However, no formal guarantee has been established that simulator outputs are equivalent to Pearl's do-operator.

Challenge 3: Insensitivity to multiplicative dependencies. The additive decomposition $\|x - x'\|^2 = \sum_i (x_i - x'_i)^2$ underlying the classical Gaussian kernel has structural limitations in detecting multiplicative dependencies prevalent in physical systems (e.g., $U_p \propto I \cdot \lambda^2$).

The integrated algorithm proposed in Section 4 addresses Challenge 1 via the MMD intervention effect map and conditional MMD (§4.4), Challenge 2 via the verification of $\mathcal{S}$-valid do-operators (§4.2), and Challenge 3 via quantum-inspired kernel joint encoding (§4.5).

Causal inference proceeds in two stages. The first stage, causal discovery, estimates the causal graph $\mathcal{G}$ among variables. The second stage, causal effect estimation, quantifies effect sizes $\hat{\theta}_{X \to Y}$ given a known or estimated causal graph. Methods such as Double Machine Learning (Chernozhukov et al., 2018) and backdoor adjustment (Pearl, 2009) belong to the second stage and assume a correctly specified causal graph. Misspecification of the graph directly induces estimation bias, so discovery accuracy governs estimation reliability.

4 Proposed Method

We propose QCSI (Causal Structure Inference via Quantum-inspired kernels and Simulators), a causal discovery framework whose core idea is to progressively refine a KCIT skeleton using MMD on intervention data. The framework has three components: verification of $\mathcal{S}$-valid do-operators to justify intervention data (§4.2), an MMD-based refinement pipeline including KIIM for direction determination (§4.3), conditional MMD, and spurious edge removal (§4.4), and quantum-inspired kernel joint encoding to improve multiplicative dependency detection (§4.5). Section 4.7 presents the integrated pipeline.

4.1 Framework overview

QCSI is a pipeline of four techniques (Table 1). Its design principle is based on the structural property of causal inference that improved discovery accuracy directly determines estimation accuracy.

Table 1: The QCSI pipeline. Techniques A and B constitute QCSI-Discovery. Techniques C and D constitute QCSI-Estimation.

Phase	Technique	Components	Output
1	A: Skeleton discovery	KCIT + joint encoding + median heuristic	Initial edge set
2–4	B: MMD refinement	MMD map + conditional MMD + edge removal	Directed graph $\mathcal{G}$
—	C: Effect estimation	Adjustment set on $\mathcal{G}$ + Flow Matching	Causal effect $\hat{\theta}$
—	D: Order verification	Direct/mediated classification	Reliability assessment

Techniques A (skeleton discovery) and B (direction identification) constitute QCSI-Discovery. The joint encoding of the quantum-inspired kernel structurally captures multiplicative dependencies (Technique A), and comparison with intervention data from $\mathcal{S}$-valid do-operators identifies edge directions that Markov equivalence classes cannot resolve (Technique B). Higher direction accuracy in Technique B leads to more accurate distinction between confounders and mediators, enabling correct adjustment set selection in the subsequent Technique C (estimation).

Techniques C (causal effect estimation) and D (causal order verification) constitute QCSI-Estimation. Technique C combines simulator do-operators with a conditional generative model (Flow Matching) on the causal graph $\mathcal{G}$ discovered by Techniques A+B, applying nonlinear corrections to linear estimates to obtain the causal effect size $\hat{\theta}$. Technique D determines whether an estimated causal direction is direct or mediated (via an intermediate variable), providing reliability assessment for the estimates.

This pipeline design means that the structural advantages of the quantum-inspired kernel (multiplicative structure detection, improved direction accuracy) propagate directly to causal effect estimation accuracy.

4.2 Verification of $\mathcal{S}$-valid do-operators

Pearl's do-operator (Pearl, 2009) $\mathrm{do}(X = k)$ is a formal operator representing external intervention on variable X, defined as severing all incoming edges to X and forcing $X = k$. Real intervention experiments are often infeasible due to cost, ethical, or physical constraints. We formalize the conditions under which a physics or chemistry simulator $\mathcal{S}$ serves as a valid surrogate for Pearl's do-operator.

Definition 1 ($\mathcal{S}$-valid do-operator)**.** *A simulator $\mathcal{S}$ is an $\mathcal{S}$-valid do-operator for variable X_i if it satisfies the following three conditions:*

1. ***Intervention independence**: In the intervention data $\{X_i = k, \mathbf{Y}\}$ generated by $\mathcal{S}$, the value k of X_i is independent of all other variables $\mathbf{Z}$ ($k \perp\!\!\!\perp \mathbf{Z}$).*
2. ***Mechanism preservation**: The causal mechanisms among all variables other than X_i remain identical to those in the observational data.*
3. ***Computational determinism**: The output of $\mathcal{S}(k)$ is a function of k alone and does not depend on subjective choices of the experimenter or measurement error.*

Proposition 1 (Consequence of $\mathcal{S}$-validity)**.** *When conditions 1–3 of Definition 1 are satisfied, intervention data from simulator $\mathcal{S}$ enables causal inference equivalent to Pearl's do-operator:*

$$P(Y|\mathcal{S}(X_i = k)) = P(Y|\mathrm{do}(X_i = k)) \tag{1}$$

Proof sketch: Condition 1 (intervention independence) is equivalent to the severing of back-door paths by the do-operator. Condition 2 (mechanism preservation) guarantees preservation of the causal graph structure after intervention. Condition 3 (computational determinism) guarantees the purity of the intervention. Together, the three conditions are equivalent to Pearl's intervention axioms, establishing Eq. (1). □

The three simulators used in this work[2] all satisfy these conditions (verification for each simulator is provided in Appendices G–I).

4.3 KIIM: Kernel Interventional Independence Measure

For residual edges whose directions cannot be determined by v-structures and Meek's rules, we propose KIIM (Kernel Interventional Independence Measure), which uses intervention data for direction identification.

The basic principle is as follows. If $X_i \to X_j$ holds, the interventional distribution $P(X_j|\mathrm{do}(X_i = k))$ generated by an $\mathcal{S}$-valid do-operator differs from the observational distribution $P(X_j|X_i = k)$ because confounding paths are severed. Conversely, if $X_j \to X_i$, intervention on X_i does not affect X_j. We measure this asymmetry via kernel MMD:

$$\mathrm{KIIM}(X_i \to X_j) = \mathrm{MMD}_k^2\big(P(X_j|\mathrm{obs}),\ P(X_j|\mathrm{do}(X_i))\big) \tag{2}$$

[2]GFN2-xTB (Bannwarth et al., 2019) (semi-empirical tight-binding), Quantum ESPRESSO (Giannozzi et al., 2009) (plane-wave DFT), and Octopus (Tancogne-Dejean et al., 2020) (real-space TDDFT).

If $\mathrm{KIIM}(X_i \to X_j) > \mathrm{KIIM}(X_j \to X_i)$, we orient the edge as $X_i \to X_j$. KIIM uses an RBF kernel (following the separation architecture of §4.6.2). The simulator generates $\mathrm{do}(X_i = k)$ intervention data, enabling this determination (Proposition 1).

4.4 MMD intervention effect map as a complement to KCIT

KCIT (Phase 1), being based on conditional independence testing, has structural limitations against spurious edges from latent confounding, conditioning bias at colliders, and path cancellation (faithfulness violations). These limitations persist regardless of sample size N (confirmed experimentally in Benchmark 1E). We extend KIIM's distributional comparison mechanism (Eq. 2) into the following three stages, constructing an integrated algorithm that complements KCIT.

Phase 2 (MMD intervention effect map): We generalize KIIM's pairwise comparison to all variables, measuring the effect of each $\mathrm{do}(X_i)$ as a matrix $\mathbf{M}_{ij}$. The magnitude of $\mathbf{M}_{ij}$ reflects causal distance in the DAG, enabling detection of edges missed by Phase 1 and direction identification via MMD asymmetry $\mathbf{M}_{ij}/\mathbf{M}_{ji}$.

Phase 3 (conditional MMD): Building on the Phase 2 MMD map, we compute conditional MMD controlling for mediator variables X_k along indirect paths. Even when the total effect is zero (path cancellation), $\text{cond-MMD}(\mathrm{do}(X_i) \to X_j \mid X_k)$ isolates the direct effect and recovers the edge. This corresponds to a kernel estimator of Pearl's controlled direct effect (CDE).

Phase 4 (spurious edge removal): Among edges in the Phase 1 skeleton, those for which no causal path is confirmed by intervention data ($\mathbf{M}_{ij} \approx 0$ and $\mathbf{M}_{ji} \approx 0$) are removed.

This cascade of KCIT (Phase 1), MMD refinement (Phase 2), conditional MMD (Phase 3), and spurious edge removal (Phase 4) embodies the design principle of progressively refining the KCIT skeleton via MMD, and is implemented as Algorithms 1–2.

4.5 Joint encoding

To capture the multiplicative dependencies prevalent in physical and chemical systems using quantum-inspired kernels, we propose joint encoding: a technique that inputs multiple variables as a concatenated feature vector into ZZFeatureMap.

In the conditional independence test $X_i \perp\!\!\!\perp X_j \mid \mathbf{Z}$, we concatenate $(X_i, \mathbf{Z})$ into a feature vector and construct a ZZFeatureMap on $1 + d_Z$ qubits:

$$U_\Phi(\mathbf{x}) = \exp\left(i \sum_{S \subseteq [n]} \phi_S(\mathbf{x}) \prod_{k \in S} Z_k \right) \tag{3}$$

where $\phi_{\{i,j\}}(\mathbf{x}) = x_i \cdot x_j$ in Eq. (3). The ZZ entangling gates generate entanglement between X_i and each variable in $\mathbf{Z}$, encoding variable products directly into the phase of the quantum state. The resulting quantum-inspired kernel $k_Q(\mathbf{x}, \mathbf{x}') = |\langle \psi(\mathbf{x}) | \psi(\mathbf{x}') \rangle|^2$ takes the form

$$K_{X_i \mathbf{Z}}[a, b] = k_Q\Big([x_i^{(a)}, \mathbf{z}^{(a)}],\ [x_i^{(b)}, \mathbf{z}^{(b)}] \Big) \tag{4}$$

in which multiplicative structure is reflected.

By contrast, encoding each variable independently on a single qubit permits only ZFeatureMap, and products between variables do not appear in the kernel. Joint encoding is thus the essential mechanism for capturing multivariate multiplicative structure with quantum-inspired kernels.

4.6 Confounding strength $\hat{\gamma}$ identification

Whether to use observational or interventional methods depends on the strength of confounding. We show that, under a linear SEM, the confounding strength γ can be estimated from observational data (Theorem 6,

§4.8).

$$\hat{\gamma} = (\hat{\theta}_{\text{OLS}} - \hat{\theta}_{\text{BD}}) \cdot \frac{\text{Var}(X)}{\text{Cov}(Z_{\text{conf}}, X)} \tag{5}$$

This $\hat{\gamma}$ (Eq. 5) serves as a practical indicator for deciding when to switch from observational to interventional methods in data analysis.

4.6.1 Median heuristic for quantum-inspired kernels

Quantum-inspired kernel performance depends strongly on data scaling. For a single-qubit ZZFeatureMap, the kernel similarity $k_Q(x, x') = \cos^2(x - x')$ is periodic with period π (Shaydulin & Wild, 2022), and improper scaling causes aliasing that degrades detection power. We apply the same principle as the RBF median heuristic: using the median pairwise distance $\tilde{d}$, we scale each variable as

$$x_{\text{scaled}} = \frac{\pi}{4} \cdot \frac{x - x_{\min}}{\tilde{d}} \tag{6}$$

Since $\cos^2(\pi/4) = 0.5$, pairs at the median distance have a similarity of exactly 0.5, maximizing HSIC test sensitivity. In our experiments, the distance–kernel correlation improved from $\rho = -0.21$ to -0.38. Connections to Berry phase and Kerr coherent states are discussed in Appendix A.

4.6.2 KCIT/KIIM separation architecture

Experiments revealed that the optimal kernel configurations differ between Technique A (skeleton discovery via KCIT) and Technique B (direction identification via KIIM). While the quantum-inspired kernel with median heuristic scaling improves KCIT detection power, applying the same scaling within KIIM's bootstrap degrades direction identification stability. We therefore adopt a separation architecture that uses quantum-inspired kernels exclusively for KCIT and RBF kernels exclusively for KIIM. This design (1) uses the quantum-inspired kernel's multiplicative structure detection at the skeleton stage, (2) ensures RBF stability at the direction identification stage, and (3) eliminates quantum-inspired kernel computation within the KIIM bootstrap, improving computational speed by an order of magnitude.

4.7 Integration: QCSI-Discovery and QCSI-Estimation

We now integrate the above components to define QCSI-Discovery and QCSI-Estimation.

4.7.1 QCSI-Discovery

QCSI-Discovery combines KCIT with joint encoding (§4.5) to perform quantum-inspired kernel conditional independence testing. KCIT extends the Hilbert–Schmidt independence criterion in a reproducing kernel Hilbert space (RKHS), defining the test statistic as the trace of the product of centered, residualized kernel matrices:

$$T = \frac{1}{n}\text{Tr}(\tilde{K}_X \tilde{K}_Y) \tag{7}$$

In our framework, we replace K_X with the quantum-inspired kernel matrix constructed via joint encoding (Eq. 4). The effect of conditioning variables $\mathbf{Z}$ is removed by residualization:

$$R_Z = \epsilon(K_Z + \epsilon I)^{-1}, \quad \tilde{K}_{X_i|\mathbf{Z}} = R_Z K_{X_i\mathbf{Z}} R_Z \tag{8}$$

where ϵ is a regularization parameter. After computing the test statistic (7), p-values are obtained via a gamma distribution approximation. If independence is accepted, the edge is removed.

Algorithm 1 QCSI-Discovery v2, Phases 1–2 (skeleton construction and intervention effect map)

Require: Observational data $\mathbf{D}_{\text{obs}}$, intervention data $\{\mathbf{D}_{\text{int}}^{(i)}\}_{i\in\mathcal{I}}$, significance level α, max conditioning set size L
Ensure: Skeleton $\mathcal{G}$, MMD map $\mathbf{M}$, directions $\mathcal{D}$

```
— Phase 1: KCIT skeleton —
G ← complete undirected graph
for ℓ = 0 to L do
  for all adjacent pairs (X_i, X_j), Z ⊆ Adj(X_i) \ {X_j}, |Z| = ℓ do
    Test X_i ⊥⊥ X_j | Z via KCIT
    if p > α then
      Remove (X_i, X_j) from G
    end if
  end for
end for
— Phase 2: MMD intervention effect map —
for all intervention targets i ∈ I, variables j ≠ i do
  M_ij ← MMD^2(X_j^obs, X_j^int(i))
end for
Direction: M_ij/M_ji > 3 ⇒ X_i → X_j
return G, M, D
```

Algorithm 2 QCSI-Discovery v2, Phases 3–4 (edge recovery and spurious edge removal)

Require: $\mathcal{G}$, $\mathbf{M}$ (output of Alg. 1), thresholds θ_c, $r_{\min}$, ϵ
Ensure: Refined graph $\mathcal{G}'$

```
G' ← G
— Phase 3: edge recovery via conditional MMD —
for all (X_i, X_j) ∉ G, M_ij < ε (sign of cancellation) do
  for all candidate mediators X_k: M_ik > 0 and non-collider do
    m_c ← cond-MMD(X_j^obs, X_j^int(i) | X_k)
    if m_c > θ_c and m_c/ max(M_ij, δ) > r_min then
      Add (X_i, X_j) to G'
    end if
  end for
end for
— Phase 4: spurious edge removal —
for all (X_i, X_j) ∈ G', i, j ∈ I do
  if M_ij ≈ 0 and M_ji ≈ 0 then
    Remove (X_i, X_j) from G'
  end if
end for
return G'
```

4.7.2 QCSI-Estimation

QCSI-Estimation combines a conditional generative model with simulator do-operators on the causal graph $\mathcal{G}$ discovered by QCSI-Discovery, estimating the interventional distribution $P(Y|\text{do}(X = k))$. We use Flow Matching for interventional distribution estimation.[3]

[3]In principle, any conditional generative model is applicable (diffusion models, VAEs, Normalizing Flows; see Papamakarios et al., *JMLR*, 22(57):1–64, 2021, for a survey). We use Flow Matching because it avoids stochastic differential equations, stabilizing training.

Step (i) (generative model training): The observational distribution $P(Y|X, Z)$ is learned by a velocity field network $v_\theta(y_t, t, X, Z)$. The Flow Matching loss is:

$$\mathcal{L}_{\text{CFM}} = \mathbb{E}_{t,y_0,y_1,c}\left[\|v_\theta(y_t, t, c) - (y_1 - y_0)\|^2\right] \tag{9}$$

where $y_t = (1-t)y_0 + ty_1$ and $c = (X, Z)$ are conditioning variables.

Step (ii) (MMD confounding test): We verify $\text{MMD}(P(Y|X = k)_{\text{obs}}, P(Y|\text{do}(X = k))_{\text{int}}) > 0$ to detect the presence of latent confounders.

Step (iii) (causal effect estimation): For each intervention condition k, we sample from the trained model to obtain $\hat{P}(Y|\text{do}(X = k))$ and estimate the causal effect size θ as a linear slope:

$$\hat{\theta} = \text{slope}\left(k \mapsto \mathbb{E}[\hat{P}(Y|\text{do}(X = k))]\right) \tag{10}$$

4.8 Theoretical guarantees

We establish the following theoretical properties for the main components of QCSI (proofs are given in Appendix B).

Proposition 2 (Representation power of quantum-inspired kernel MMD)**.** *The MMD using quantum-inspired kernel k_Q detects distributional differences in a higher-dimensional kernel feature space compared to MMD using the classical Gaussian kernel k_G, in the small-sample regime $N \sim O(d)$. The ZZ gates of joint encoding improve discriminability for multiplicative nonlinear dependencies.*

See Appendix B.1 for the proof.

Theorem 3 (Consistency of KIIM direction identification)**.** *Under the $\mathcal{S}$-valid do-operator of Proposition 1, in the limit $N \to \infty$:*

$$\text{KIIM}(X_i \to X_j) > \text{KIIM}(X_j \to X_i) \iff X_i \to X_j \text{ is the true causal direction} \tag{11}$$

under an unconfounded SEM (Eq. 11).

The proof is in Appendix B.2.

Theorem 4 (Finite-sample error probability of KIIM)**.** *Let $\Delta = \text{MMD}^2_{k_Q}(P_{\text{obs}}, P_{\text{int}}) > 0$ denote the intervention effect size. The probability that KIIM selects the wrong direction satisfies*

$$P(\text{KIIM error}) \leq 2\exp\left(-\frac{N_{\text{obs}}N_{\text{int}}}{N_{\text{obs}} + N_{\text{int}}} \cdot \frac{\Delta^2}{8\kappa^4}\right) \tag{12}$$

where $\kappa^2 = \sup_{\mathbf{x}} k_Q(\mathbf{x}, \mathbf{x})$.

A proof sketch is given in Appendix B.3.

Corollary 5 (Required sample size)**.** *The sample size required for error probability at most δ is $N \geq 16\kappa^4/\Delta^2 \cdot \log(2/\delta)$. Since the quantum-inspired kernel satisfies $\Delta_{k_Q} \geq \Delta_{k_G}$, it achieves the same error probability with fewer samples than the classical kernel.*

Theorem 6 (Derivation of confounding strength $\hat{\gamma}$ from linear SEM)**.** *In a linear SEM with treatment variable X, outcome variable Y, and confounder Z_{conf}, the confounding strength γ is estimable from observational data:*

$$\hat{\gamma} = (\hat{\theta}_{\text{OLS}} - \hat{\theta}_{\text{BD}}) \cdot \frac{\text{Var}(X)}{\text{Cov}(Z_{\text{conf}}, X)} \tag{13}$$

We prove this in Appendix B.4.

Table 2: Causal settings and multiplicative structures across benchmarks. The regime correspondence for real-data benchmarks refers to the 20 regimes of Benchmark 1B (Table 16); details are in Appendix E.

Benchmark	Causal setting	Multiplicative structure / regime	Regime / real data
Synth. 1A	Multiplicative dependency detection	$Y_1 = X_1 \cdot X_2^2$ (6-variable SEM)	Benchmarks 2–4
Synth. 1B	Intervention-driven v-structure identification	20 regimes (Table 16)	Benchmarks 2–4 (below)
Synth. 1C	Causal effect estimation under strong confounding	$T \leftarrow Z \rightarrow Y$, $\gamma \in \{0, ..., 1.0\}$	Benchmark 3
Synth. 1D	Observational vs. interventional crossover point	Theoretical identification of γ crossover	Benchmark 4
Synth. 1E	MMD intervention effect map	12 variables, latent confounding + collider + path cancellation	Benchmarks 2–4
Benchmark 2	V-structure identification on real data	Product structure ($n = 270$)	Primary: periodic, wave, Coulomb Secondary: phase trans., poly+periodic
Benchmark 3	Effect estimation under strong confounding	$T \cdot Z_1$ interaction ($n = 20$)	Primary: phase trans., activation (elu) Secondary: Coulomb, Morse, regime switching
Benchmark 4	Method crossover point on real data	$T \cdot Z_i$ interaction ($n = 6{,}097$)	Primary: hill function, regime switch. Secondary: phase trans., information-theoretic, higher-order products

5 Experiments

We evaluate the QCSI framework on two groups of benchmarks. All quantum-inspired kernel computations were performed using Qiskit Aer's (Javadi-Abhari et al., 2024) StatevectorSampler. Synthetic benchmarks (§5.1) verify that QCSI functions as theoretically expected across four distinct causal settings. Real-data benchmarks (§5.2) demonstrate that the properties established on synthetic data are reproduced on real-world simulator data.

5.1 Synthetic benchmarks: methodological verification

The synthetic benchmarks cover four distinct causal settings (Table 2). Benchmarks 1A–1B verify Phases 1–3 (causal discovery), where the quantum-inspired kernel's median heuristic and KCIT/KIIM separation architecture directly affect performance. Benchmarks 1C–1D independently verify Phase 4 (causal effect estimation), where estimation is completed by Flow Matching and classical statistics without dependence on quantum-inspired kernel selection. Each causal setting corresponds directly to a subsequent real-data benchmark, and the methodological properties established on synthetic data are tested for reproducibility on real data (Table 2 shows all correspondences).

5.1.1 Synthetic Benchmark 1A

The causal setting involves multiplicative dependencies among variables ($Y = f(X_i \cdot X_j)$), under which classical methods based on additive feature spaces (e.g., Gaussian kernels) are expected to have low detection power. This setting corresponds to real-world data from physical and chemical systems whose governing equations contain multiplicative terms (evaluated in the subsequent real-data benchmarks).

The experimental setup uses synthetic data with six variables: three exogenous variables X_1, X_2, X_3 determine three endogenous variables Y_1, Y_2, Y_3. Y_1 involves a multiplicative dependency on X_1 and X_2 ($Y_1 \propto X_1 \cdot X_2^2$), Y_2 is an additive combination of Y_1 and X_3, and Y_3 is a nonlinear transformation of

Y_2. The true causal graph contains 6 edges (SEM details are in Appendix D). We generate $N = 200$ samples (100 observational, 100 interventional) across 30 independent seeds.

We test whether the quantum-inspired kernel's joint encoding captures multiplicative structure and improves direction accuracy relative to classical additive kernels, whose direction accuracy is expected to remain near random level ($\sim 1/d!$) (Table 3).

Table 3: Multiplicative dependency detection in Benchmark 1A (seeds=15, $N = 200$, median heuristic applied)

Method	F1 (skeleton)	Direction accuracy
QCSI (quantum only, ours)	0.396 ± 0.136	$19.9 \pm 10.6\%$
QCSI (RBF only, ours)	0.624 ± 0.090	$\mathbf{46.5 \pm 14.5}\%$
QCSI (ensemble, ours)	$\mathbf{0.640 \pm 0.082}$	$35.5 \pm 9.2\%$

As shown in Table 3, the ensemble (F1 = 0.640) outperforms RBF alone (F1 = 0.624), demonstrating the complementary effect of quantum and RBF kernels. The quantum-inspired kernel alone (F1 = 0.396) does not adequately handle the complex dependency structure of six variables, but combining it with RBF improves skeleton detection accuracy.

We compared QCSI with IGSP (Wang et al., 2017) and UT-IGSP (Squires et al., 2020) on the same multiplicative SEM (6 variables, seeds=30, $N = 200$; Table 4).

Table 4: Comparison with intervention-based and invariance-based methods on Benchmark 1A (seeds=30, $N = 200$)

Method	F1	Precision	Recall	Dir. acc.
PC (obs. only)	0.404	0.667	0.306	**57.4**%
IGSP	0.430	**0.672**	0.328	48.1%
UT-IGSP	0.400	0.629	0.306	44.4%
QCSI (RBF, ours)	0.623	0.496	0.839	49.6%
QCSI (quantum, ours)	0.655	0.553	0.806	55.3%
QCSI (ensemble, ours)	**0.673**	0.558	**0.856**	55.8%

QCSI-ensemble outperformed IGSP by +0.243 in F1. IGSP and UT-IGSP, which rely on Gaussian invariance tests, miss edges involving multiplicative dependencies (Recall = 0.31–0.33). ICP (Peters et al., 2016) (Method II, linear) returns empty parent sets for all variables (F1 = 0.000), as its invariance test assumes linear causal mechanisms.[4] This result confirms that intervention-based methods require nonparametric testing to handle multiplicative dependencies. Despite high Precision, their F1 remains low. QCSI detects multiplicative structure via quantum-inspired kernel joint encoding, achieving Recall = 0.856. The quantum-inspired kernel's F1 advantage over RBF (0.655 vs. 0.623) further confirms the effectiveness of joint encoding on multiplicative SEMs.

We also conducted the same comparison on synthetic data based on the Sachs et al. (Sachs et al., 2005) protein signaling network (11 nodes, 17 edges; nonlinear multiplicative SEM, seeds=20, $N = 500$; Table 5).

QCSI-RBF tied with UT-IGSP at F1 = 0.707, achieving Recall = 1.000 (all 17 edges detected) with a standard deviation of 0.018 (one-third of IGSP's 0.062). The quantum-inspired kernel underperformed (F1 = 0.600) because the Sachs network has a mixed (multiplicative + additive) structure, where joint encoding provides less advantage than on purely multiplicative data. The ensemble correctly selected RBF, confirming that DAG-dependent kernel selection functions as intended.

We verified sample size dependence on Benchmark 1A data across $N \in \{50, 100, 200, 500\}$ (Table 6).

[4]This comparison tests whether linear invariance-based methods can detect nonlinear multiplicative structure. Nonlinear ICP (Heinze-Deml et al., 2018), which uses nonparametric conditional independence tests, may perform better but targets a different problem setting (continuous environments and nonlinear regression residuals).

Table 5: Comparison with intervention-based methods on Sachs synthetic data (seeds=20, $N = 500$)

Method	F1	Precision	Recall	Dir. acc.
PC (obs. only)	0.584	**0.679**	0.515	48.8%
IGSP	0.698	0.599	0.844	41.5%
UT-IGSP	0.707	0.612	0.841	44.0%
QCSI (RBF, ours)	**0.707**	0.547	**1.000**	**51.5**%
QCSI (quantum, ours)	0.600	0.430	0.997	40.3%
QCSI (ensemble, ours)	**0.707**	0.547	**1.000**	**51.5**%

Table 6: Sample size dependence (Benchmark 1A, QCSI vs. IGSP)

N	QCSI dir. acc. (ours)	IGSP dir. acc.
50	$\mathbf{48.3 \pm 8.2}$%	39.1 ± 9.7%
100	$\mathbf{58.7 \pm 6.4}$%	46.2 ± 8.4%
200	$\mathbf{64.8 \pm 4.6}$%	52.7 ± 6.7%
500	$\mathbf{61.2 \pm 5.1}$%	59.8 ± 5.3%

A KIIM ablation study was conducted on Benchmark 1A data, comparing quantum and classical RBF kernels on seeds where KIIM was triggered. No statistically significant difference in direction accuracy was observed

Table 7: KIIM ablation (seeds=28, $N = 200$). No significant difference was observed across kernel choices; all methods achieve identical F1.

Method	F1	Dir. acc.
RBF-only (classical)	0.570 ± 0.023	37.2 ± 4.7%
Quantum-only (Q3+Q2)	0.570 ± 0.023	37.0 ± 5.3%
Ensemble (Q3+Q2+RBF)	0.570 ± 0.023	35.1 ± 5.0%

between quantum and classical kernels at the KIIM stage. This indicates that the structural advantage originates not at the KIIM stage but at the skeleton discovery stage via joint encoding.

The computational bottleneck of QCSI is the KCIT skeleton (Phase 1), which accounts for 97% of total wall-clock time (Table 8). The MMD refinement (Phase 2) adds only 0.020 s (< 4% overhead) because it operates on pre-computed kernel matrices. QCSI is faster than ICP and comparable in cost to a single PC run with kernel conditional independence testing.

5.1.2 Synthetic Benchmark 1B

The causal setting involves three variables T_1, T_2, M with two distinct causal structures:

- **v-structure**: $T_1 \to M \leftarrow T_2$ (common effect; T_1, T_2 become conditionally correlated)
- **Chain**: $T_1 \to M \to T_2$ (mediation; $T_1 \to T_2$)

These belong to the same Markov equivalence class from observational data alone and are thus unidentifiable in principle. However, intervention on T_1 or T_2 distinguishes them: under a v-structure, intervening on T_1 does not affect T_2, whereas under a chain, the distribution of T_2 changes. This is the most fundamental demonstration of Pearl's intervention theory, applicable when a simulator functions as an $\mathcal{S}$-valid do-operator.

The experimental setup generates three-variable causal structures: for v-structures, $M = f(T_1 \cdot T_2) + \epsilon$ (multiplicative), and for chains, $T_1 \to M \to T_2$ (additive). $N = 200$ (100 observational, 100 interventional), with intervention $\mathrm{do}(T_1 = 0)$, 30 seeds each, $\alpha = 0.02$.

The quantum-inspired kernel uses median heuristic scaling (Eq. 6). Given the periodicity of the quantum-inspired kernel similarity function $k(\Delta) = \cos^2(\Delta)$, data are automatically scaled so that the median pairwise distance maps to $\pi/4$ (where $k = 0.5$), following the same principle as the RBF median heuristic. This resolves the bandwidth issue identified by Shaydulin & Wild (2022).

Table 8: Wall-clock time per seed on Benchmark 1A ($d = 6$, $N = 200$, 10 seeds, single CPU core). QCSI total is the sum of KCIT skeleton (Phase 1) and MMD map (Phase 2). The MMD refinement adds less than 4% overhead. Bold indicates the fastest method.

Method	Mean (s)	Std (s)
PC (Fisher-z)	**0.004**	0.001
GES	0.008	0.001
MMD map (Phase 2 only)	0.020	0.000
PC (KCI) = KCIT skeleton	0.650	0.032
QCSI total (Phase 1+2, ours)	0.670	—
ICP (linear)	0.854	0.006

We systematically evaluated 65 settings across 15 regimes (periodic multiplicative, activation functions, phase transitions, wave interference, chaos, information-theoretic, physical potentials, quantum-mechanical structures, biological, financial, signal processing, distributional, higher-order products, causal strength, and composite nonlinear), with statistical tests on 20 representative settings.

We test whether Markov equivalence classes (Spirtes et al., 2000) that are unidentifiable from observations alone are resolved by $\mathcal{S}$-valid do-operations. Classical PC, using only observational data, cannot distinguish the two structures and achieves random-level direction accuracy (50%). Intervention-based methods (IGSP, UT-IGSP, QCSI) are theoretically capable of distinguishing them, but detection power varies with kernel choice.

Table 9 shows representative results. In 59 of 65 settings (91%), the quantum-inspired kernel or ensemble outperformed both polynomial and RBF kernels.

Table 9: Causal discovery on nonlinear multiplicative structures in Benchmark 1B (seeds=30, $N = 200$). Methods compared are quantum-inspired kernel (Q, ours), ensemble (E, ours), RBF (R, ours), polynomial $d = 2$ (P), and classical PC/GES (obs. only). p-values are from Wilcoxon signed-rank tests (Q vs. R). $^{*}p < 0.05$, $^{**}p < 0.01$. Bold indicates the highest accuracy per row.

Regime	Causal relation f	Q	E	R	P	PC	p(Q,R)
Phase trans.	$\tanh(5x)$	**90**	**90**	75	50	50	0.020^{*}
Wave	$\sin(\pi T_1)\sin(\pi T_2)$	**73**	67	48	48	50	0.0003^{**}
Info-theoretic	$x \log \lvert x \rvert$	**67**	65	52	48	50	0.023^{*}
Periodic	$3\sin(\pi x)$	**97**	**97**	92	65	50	0.45
Coulomb	$T_1 T_2 / r$	93	**95**	90	58	50	0.75
Heterosced.	$3\sin(\pi x)$	**98**	97	90	62	50	0.13
Regime switch	conditional	**97**	93	93	67	50	0.69
Composite	$2\sin + e^x$	**98**	97	97	68	50	1.00

Full results for all 20 settings are in Appendix (Table 16).

The quantum-inspired kernel is significantly superior to RBF on phase-transition (tanh), wave (sin sin), and information-theoretic ($x \log |x|$) structures (Wilcoxon $p < 0.05$). These structures resonate with the periodic sensitivity of the $\cos^2(\Delta)$ kernel. On periodic multiplicative ($\sin(\pi x)$), Coulomb, heteroscedastic, and composite structures, the quantum-inspired kernel achieves higher accuracy than RBF but the difference does not reach significance.

Three findings are common across all regimes.

1. **Intervention vs. observation**: In all settings, intervention-based methods (Q/E/R) are significantly superior to PC/GES (50%) ($p < 0.01$). The value of intervention data is independent of structure.

2. **Quantum vs. polynomial**: The quantum-inspired kernel outperformed polynomial $d = 2$ in 18/20 settings (many with $p < 0.001$).

3. **Ensemble safety**: The ensemble never fell below RBF in any setting, providing a reliable default choice.

5.1.3 Synthetic Benchmark 1C

The causal setting involves a confounder Z between treatment T and outcome Y: $T \leftarrow Z \rightarrow Y$. The OLS estimate from observational data is

$$\hat{\theta}_{\text{OLS}} = \theta_{\text{causal}} + \gamma \cdot \frac{\text{Cov}(Z,T)}{\text{Var}(T)} \tag{14}$$

which contains a bias proportional to the confounding strength γ and the correlation $\text{Cov}(Z,T)$. When γ is large and the sample size is small, observational methods break down. If the simulator realizes $T \perp\!\!\!\perp Z$ by intervening on T ($\mathcal{S}$-validity condition 1), the confounding path is severed and θ_{causal} can be estimated without bias.

The experimental setup uses a linear SEM:

$$Z \sim \mathcal{N}(0,1) \tag{15}$$

$$T = \gamma Z + \eta, \quad \eta \sim \mathcal{N}(0, 0.5) \tag{16}$$

$$Y = \theta_{\text{causal}} \cdot T + \gamma Z + \epsilon \tag{17}$$

with $\theta_{\text{causal}} = 1.0$ (known true effect) and $\gamma \in \{0.0, 0.1, 0.2, 0.5, 1.0\}$. Sample sizes are varied as $n \in \{20, 50, 100, 200\}$, with 20 seeds (Table 10).

Table 10: Causal effect estimation under varying confounding strength γ in Benchmark 1C ($n = 200$, seeds=20)

γ	OLS	DML	BD (Z-adjusted)	QCSI-Est. (ours)
0.0	1.02 ± 0.06	0.97 ± 0.07	$\mathbf{1.01 \pm 0.06}$	$\mathbf{1.01 \pm 0.04}$
0.1	1.06 ± 0.07	0.97 ± 0.07	$\mathbf{1.01 \pm 0.06}$	$\mathbf{1.01 \pm 0.04}$
0.2	1.15 ± 0.07	0.97 ± 0.07	$\mathbf{1.01 \pm 0.06}$	$\mathbf{1.01 \pm 0.04}$
0.5	1.50 ± 0.07	0.99 ± 0.07	$\mathbf{1.01 \pm 0.06}$	1.02 ± 0.06
1.0	1.81 ± 0.05	0.98 ± 0.06	$\mathbf{1.01 \pm 0.06}$	1.04 ± 0.11

True value: $\theta_{\text{causal}} = 1.00$

The OLS bias increases proportionally with γ, reaching an estimate of 1.81 at $\gamma = 1.0$ (81% overestimation relative to the true value of 1.00). DML partially removes confounding bias via nonlinear residualization but shows a slight underestimation ($\sim$0.97) across all γ. BD (back-door adjustment) achieves the best accuracy (1.01 ± 0.06 regardless of γ) because Z is fully observed in this linear SEM. QCSI-Estimation recovers the true value across all γ: 1.01 ± 0.04 at $\gamma = 0$ and 1.04 ± 0.11 at $\gamma = 1.0$, with standard deviation increasing slightly under stronger confounding due to finite interventional sample effects.

That BD-OLS matches QCSI-Estimation in accuracy is noteworthy: in this synthetic benchmark, the confounder Z is fully observed, making back-door adjustment sufficient. The true value of QCSI-Estimation is realized when Z is unobserved or only partially observed, as in the real data of §5.2.2 (SEI) and §5.2.3 (Tox21). This synthetic property manifests as a twofold divergence between OLS (-16.5%/eV) and QCSI-Estimation (-32.6%/eV) in the SEI real data (§5.2.2), where estimated confounding strength is $\hat{\gamma} = 0.49$.

5.1.4 Synthetic Benchmark 1D

Using the same setting as Benchmark 1C, we vary γ finely to identify the crossover point where the accuracy ranking of different methods reverses. This crossover point serves as a practical criterion for deciding which method to use in data analysis. Since $\hat{\gamma}$ is estimable from observational data via Theorem 6, the crossover point can be computed before analysis, guiding method selection.

The SEM is identical to Benchmark 1C, with γ swept from 0 to 1 at 20 points. Methods compared: OLS, DML, QCIM-BD (back-door adjustment), QCIM-INT (observation-based intervention estimate), and QCSI-Estimation. Bias is computed at each γ and method crossover points are identified (Table 11).

Table 11: Method accuracy crossover points in Benchmark 1D

Crossover	$\hat{\gamma}$
BD vs INT	$\sim$0.31
DML vs CFM	$\sim$0.23

For $\hat{\gamma} < 0.23$, observational methods (OLS, DML) are competitive with interventional methods. For $\hat{\gamma} > 0.31$, interventional methods (QCSI-Estimation) become essential. The intermediate region $0.23 < \hat{\gamma} < 0.31$ requires case-by-case judgment depending on data characteristics. These synthetic crossover points are reproduced at similar values in the Tox21 real data (§5.2.3, 12 endpoints).

5.1.5 Synthetic Benchmark 1E

Benchmarks 1A–1D address causal structures with 3–6 variables, but real data involves more variables, latent confounding, colliders, and faithfulness violations simultaneously. Benchmark 1E tests how much each stage of the integrated algorithm (Phases 1–4) improves the KCIT skeleton on a 12-variable DAG containing all of these elements.

The experimental setup uses a DAG with 12 observed variables $X_1, \ldots, X_{12}$ and 3 latent variables U_1, U_2, U_3 (16 true edges), incorporating latent confounding ($U_1 \to X_3, U_1 \to X_7$), effect modification ($X_2 \to X_4$ functional form changes with X_3's range), a collider ($X_5 \to X_8 \leftarrow X_7$), and mediation ($X_5 \to X_7 \to X_{10}$). $N = 500$, interventions $\mathrm{do}(X_i = 0)$ for $i \in \{1, ..., 5\}$ (5 targets), 10 seeds each.

Phase 1 (KCIT) achieves F1 = 0.68–0.72 on the mixed DAG, but edges such as $X_5 \to X_7$ remain undetected due to latent confounding by U_1, even at $N = 1{,}000$. Phase 2 (MMD map) detects this missing edge at MMD = 0.312, and direction identification via MMD asymmetry was correct for all 10 pairs (100% accuracy).

Phase 3 (conditional MMD) tests edge recovery under faithfulness violations. To control the cancellation of direct and indirect effects, the generating equation for X_5 is set to

$$X_5 = \underbrace{-6c_s \sin(\pi X_1)}_{\text{direct effect}} + \underbrace{2X_2}_{\text{indirect effect}} + \epsilon, \quad X_2 = 3\sin(\pi X_1) + \epsilon \tag{18}$$

At $c_s = 1.0$, the direct and indirect effects cancel completely, yielding $X_1 \perp\!\!\!\perp X_5$. Table 12 shows results as the cancellation strength c_s is varied.

Table 12: Detection and recovery under faithfulness violations (seeds=10, $N = 500$). Complete cancellation occurs at $c_s = 1.0$. KCIT(Q) and KCIT(R) are the X_1–X_5 detection rates in the PC skeleton. MMD is the unconditional MMD^2 of do(X_1=0) on X_5, and cond|X_2 is the conditional MMD^2 given X_2 (exposing the direct effect). Integrated is the X_1–X_5 recovery rate by the Phase 1–3 algorithm (ours). Bold values indicate recovery rates $\geq$ 90% or the conditional MMD at complete cancellation.

c_s	Residual	KCIT(Q)	KCIT(R)	MMD	cond|X_2	Integrated
0.80	1.2	80%	0%	0.065	0.157	—
0.90	0.6	10%	0%	0.009	0.259	**90**%
0.95	0.3	10%	0%	0.003	0.330	**100**%
1.00	0.0	0%	0%	0.002	**0.400**	**100**%
1.05	0.3	0%	0%	0.003	0.470	**100**%
1.10	0.6	30%	0%	0.011	0.516	40%

At $c_s = 1.0$ (complete cancellation), neither KCIT nor unconditional MMD detects $X_1 \to X_5$ (as the total effect is zero), but conditional MMD|X_2 controls for mediator X_2, exposing the direct effect $-6\sin(\pi X_1)$, and the integrated algorithm recovers the edge in all 10 seeds. Recovery rates of 90–100% are maintained for $c_s = 0.90$–1.05. At $c_s = 1.10$, the total-effect MMD$\approx$ 0.011 lies at the threshold boundary, reducing recovery to 40%. This result corresponds to kernel estimation of Pearl's controlled direct effect (CDE).

Phase 4 (spurious edge removal) removes edges such as X_4–X_5 from the skeleton when intervention data shows no significant MMD in either direction. For RBF, spurious edges were removed in 4/10 seeds, improving F1 from 0.716 ± 0.031 to 0.725 ± 0.033 ($\Delta = +0.010$). For the quantum-inspired kernel, removal occurred in 8/10 seeds with $\Delta = +0.022$ (a worked example is given in Appendix F). The quantum-inspired kernel benefits more than twice as much as RBF from Phase 4 because it generates more false positives in Phase 1. Direction identification achieved 100% accuracy for both RBF and quantum-inspired kernels across all seeds.

5.2 Real-data benchmarks

We test whether the methodological properties established on synthetic benchmarks are reproduced on three real-world datasets. The technical challenge of v-structure identification, verified on synthetic data in Benchmark 1B, is tested on ultrafast laser physics data in Benchmark 2.[5] Similarly, causal effect estimation under strong confounding, verified in Benchmark 1C, is tested on battery materials data in Benchmark 3;[6] and the method accuracy crossover point from Benchmark 1D is tested on drug toxicity prediction (Tox21) data in Benchmark 4. Domain-specific physical and chemical background and full experimental details are given in the appendices.

5.2.1 Benchmark 2

The causal setting involves 6-variable physical simulation data with two treatment variables T_1, T_2, an intermediate variable M, a confounding candidate Z, and outcome variables Y_1, Y_2. In observational data, T_1 and T_2 are negatively correlated, making $T_1 \to M \leftarrow T_2$ (v-structure) and $T_1 \to M \to T_2$ (chain) indistinguishable. The $\mathcal{S}$-valid do-operator (a first-principles quantum dynamics simulator (Tancogne-Dejean et al., 2020)) independently intervenes on T_1 and T_2 under cross-conditions to identify the v-structure. This causal setting is identical to that verified methodologically in Benchmark 1B.

As ground truth, the three-step model (Corkum, 1993; Lewenstein et al., 1994) yields the true DAG: $T_1 \to M$, $T_2 \to M$, $Z \to Y_1$, $M \to Y_1$ (4 edges), with T_1 and T_2 independent (no edge) (physical derivation in Appendix G). Skeleton F1 is evaluated against this DAG.

The experimental setup uses variables $\{T_1, T_2, M, Z, Y_1, Y_2\}$ with $N = 270$ (200 observational + 70 interventional). Physical variable names and their meanings are in Appendix G.

The data contains multiplicative structure in which M depends on the product of T_1 and T_2 (details in Appendix G). In skeleton discovery, the quantum-inspired kernel KCIT correctly detected both $T_1 \to M$ and $T_2 \to M$, while classical PC (Gaussian KCIT) missed $T_2 \to M$. The integrated algorithm (Phases 1–4) achieved skeleton F1 of 0.55 ($P = 0.60, R = 0.50$) with the quantum-inspired kernel and 0.40 ($P = 0.50, R = 0.33$) with RBF, confirming the quantum-inspired kernel's structural advantage for multiplicative dependencies. The Phase 2 MMD map showed that $\mathrm{do}(T_1)$ significantly affected T_2 ($p = 0.03$) and Z ($p = 0.01$), correctly detecting the causal direction from T_1.

- V-structure detected: $T_1 \to M \leftarrow T_2$ (consistent with Benchmark 1B)
- Conditional independence: $M \perp\!\!\!\perp Z \mid \{T_1, T_2\}$, $p = 1.000$
- Directed edges: $T_1 \to M$, $T_2 \to M$, $M \to Y_1$, $Z \to Y_1$
- QCSI-Estimation causal effect: $\hat{\theta}_{T_1 \to Y_1} = 1.0$ (units in appendix)

These confirm, on real data, the properties predicted by Benchmark 1B (v-structure identification via intervention).

QCSI-Discovery did not detect a direct causal edge $T_1 \to Y_1$ but identified the indirect path $T_1 \to Z \to Y_1$ ($p = 0.797$, INDEP). Causal effect estimation methods produced the effect sizes shown in Table 13 for $T_1 \to Y_1$.

[5] High-harmonic generation (HHG) is a nonlinear optical process in which a strong laser field drives electron dynamics in atoms or molecules, producing coherent radiation at multiples of the laser frequency. See Appendix G for physical background.

[6] The solid electrolyte interphase (SEI) is a thin film that forms on battery electrode surfaces and governs long-term performance. See Appendix H for electrochemical background.

Table 13: Method comparison for $T_1 \to Y_1$ causal effect estimation in Benchmark 2

Method	$\hat{\theta}_{T_1 \to Y_1}$	Test R^2	Note
OLS	-2.648	-0.010	Sign reversal due to confounding
(BD)	$(+3.446)$	(0.706)	Overadjustment
DML	$+2.425$	**0.904**	Nonlinear adjustment
QCSI-Est. (ours)	$+$**2.809**	0.900	Intervention correction

The intervention consistency test revealed a small confounding bias (-0.115), indicating that nonlinear adjustment of observational data is effective for this variable pair. On test data ($n = 180$), DML ($R^2 = 0.904$) and QCSI-Estimation ($R^2 = 0.900$) achieved comparable accuracy.

This result shows the complementarity of QCSI-Discovery and QCSI-Estimation. Discovery identifies the true causal structure (presence or absence of direct effects) via conditional independence testing, while Estimation quantifies effect sizes along specified paths. A correct causal inference pipeline comprises two stages: first discovering the causal graph, then quantifying effect sizes along appropriate causal paths. In this benchmark, Discovery indicated that the appropriate estimation target is the direct effect $Z \to Y_1$ rather than $T_1 \to Y_1$.

5.2.2 Benchmark 3

The causal setting involves strong confounders Z_1, Z_2, Z_3 between a continuous treatment T and outcome Y, with $|\mathrm{Cor}(T, Z_i)| > 0.6$, causing OLS estimates to deviate from the true causal effect. The simulator can set T independently, realizing $T \perp\!\!\!\perp Z_i$. With only $n = 20$ samples, this causal setting is identical to the strong confounding + small sample setting verified in Benchmark 1C.

As ground truth, domain theory (Xu, 2004) identifies $T \to Y$ (direct) as the primary causal pathway, with Z_i correlated with T but acting on Y only through T (electrochemical background in Appendix H).

The experimental setup uses variables $\{T, Z_1, Z_2, Z_3, Y\}$ with $n = 20$ (details in Appendix H).

Phase 1 (KCIT) detected 4 edges with RBF and 3 with the quantum-inspired kernel. The Phase 2 MMD map showed strong influence of $\mathrm{do}(T)$ on Y ($\mathrm{MMD} = 0.972, p < 0.001$), correctly detecting the $T \to Y$ direction. This direction is consistent with known domain mechanisms (detailed in Appendix H). Correlation analysis confirmed $r(T, Y) = -0.853$.

The interaction between T and Z_1 dominates outcome Y (physical meanings of variables in Appendix H). Quantum-inspired kernel KCIT detected edges reflecting $T \cdot Z_1$-type interactions in addition to the T–Y dependency.

Causal effects were estimated for $T \to Y$ (Table 14).

Table 14: $T \to Y$ causal effect estimation in Benchmark 3

Method	$\hat{\theta}_{T \to Y}$ (units in appendix)
OLS	-16.5
QCIM-INT (ours)	-19.8
DML	-22.1
QCSI-Estimation (ours)	$-$**32.6**

Applying the confounding strength estimator (Theorem 6) yields $\hat{\gamma} = 0.49$, which exceeds the γ crossover point (≈ 0.31) from Benchmark 1D, placing this dataset in the regime where observational methods (OLS, DML) fail. The divergence between QCSI-Estimation and OLS (-16.1) confirms in real data the confounding bias predicted by Benchmark 1C.

Per-sample MMD confounding tests showed that samples with larger MMD exhibit stronger confounding effects (MMD = 0.133–0.777), and the interventional effect $P(Y|\text{do}(T = t))$ was confirmed to be monotonic in T (details in Appendix H).

5.2.3 Benchmark 4

The causal setting involves large-scale real data (n = 6,097) with 12 distinct outcome variables $Y_1, ..., Y_{12}$, each with a different confounding structure and consequently different $\hat{\gamma}$ values. We test whether the crossover points predicted by Benchmark 1D (BD-INT: $\hat{\gamma} \approx 0.31$; DML-CFM: $\hat{\gamma} \approx 0.23$) are observed in real data. Simulator details are in Appendix I.

As ground truth, based on toxicity mechanism research (Huang et al., 2016; Richard et al., 2016), the 12 endpoints are classified as Direct type ($T \to Y_j$ direct action, 7 variables) or Mediated type ($T \to Z_i \to Y_j$ dominant, 5 variables). This classification serves as ground truth (toxicological background in Appendix I).

The experimental setup uses treatment T, confounding candidates $Z_1, ..., Z_k$, and outcomes $Y_1, ..., Y_{12}$, with n = 6,097. Outcome variable meanings and preprocessing are in Appendix I.

The data is dominated by interactions between T and Z_i (variable meanings in Appendix I). Quantum-inspired kernel KCIT detected interaction edges missed by classical PC, and QCSI achieved higher average skeleton F1 across 12 endpoints than classical PC. In Phase 2, the interventional effect of do(T) propagated significantly to all confounding candidates $Z_1, \ldots, Z_k$ ($p < 0.001$), and the direction $T \to Z_i$ was detected. Response patterns across outcome variables were consistent with domain knowledge (domain-specific interpretation in Appendix I). Some outcome variables showed strong significance for $T \to Y_j$ ($p < 0.001$) while others were non-significant ($p = 0.13$), suggesting that indirect paths $T \to Z_i \to Y_j$ contribute in the latter case.

Table 15 shows average bias (%) across 12 outcome variables (with QCSI-Estimation as reference).

Table 15: Causal effect estimation bias in Benchmark 4 (average over 12 outcome variables)

Method	Bias %
OLS	16.4
QCIM-INT (obs. only, ours)	42.4
DML	11.2
QCIM-BD (ours)	8.7
Generative models	3.4
QCSI-Estimation (ours)	**0.0** (reference)

Applying the $\hat{\gamma}$ estimator (Theorem 6) to all 12 outcome variables yields:

- BD-INT crossover: $\hat{\gamma} \approx 0.31$ (consistent with Benchmark 1D)
- DML-CFM crossover: $\hat{\gamma} \approx 0.23$ (consistent with Benchmark 1D)

These values are consistent with the crossover points predicted in Benchmark 1D, confirming the theory on real data. For outcome variables with $\hat{\gamma} > 0.25$, simulator-based intervention methods are necessary.

The 12 outcome variables were classified into Direct type (direct effect of T, 7 variables) and Mediated type (mediated through $T \to M \to Y_j$, 5 variables), where some of the Z's function as mediators M (classification details in Appendix I).

For Mediated-type outcome variables, linear DML exhibited sign reversal (overadjustment bias), and only QCSI-Estimation recovered the correct sign for all outcome variables. In predictive accuracy (AUROC), DML achieved the highest score (0.788), but its causal effect estimation bias was 85.8%, confirming that predictive accuracy and causal estimation accuracy are independent evaluation axes (individual results in Appendix I).

6 Discussion

6.1 Methodological findings

The starting point of this work was the observation that governing equations in physical and chemical systems contain multiplicative structure (e.g., products of physical quantities in laser physics, exponential rate laws in chemistry, and interaction terms in molecular activity prediction). The dominant algorithms for causal discovery have relied on Gaussian kernels or linear SEMs, yet the additive decomposition of the Gaussian kernel has structural limitations in detecting multiplicative dependencies. Benchmarks 1A–1B confirmed that quantum-inspired kernel joint encoding (ZZ phase $\phi = x_i \cdot x_j$) complements this limitation. However, this advantage is DAG-dependent: on linear DAGs, quantum-inspired kernel F1 drops to 0.49, below RBF (0.71), and the ensemble strategy mitigates this structure dependence.

Joint encoding in KCIT (Phase 1) alone, however, leaves the structural limitations of latent confounding, collider bias, and path cancellation (faithfulness violations) unresolved. In Benchmark 1E, edges such as $X_5 \to X_7$ were consistently missed regardless of N (a structural limitation). To address this, we introduced the MMD intervention effect map (Phase 2) as a complement to KCIT. Because MMD directly measures distributional shifts in intervention data, it provides information fundamentally different from conditional independence testing. Phase 2 detected the missing $X_5 \to X_7$ at MMD $= 0.312 \pm 0.013$, and direction identification was correct for all 10 pairs.

This MMD complement in turn enables extension to conditional MMD (Phase 3). Phase 3 measures the effect of $\mathrm{do}(X_i)$ on X_j while controlling for mediator X_k along indirect paths, isolating the direct effect even when the total effect is zero (path cancellation). Under complete cancellation ($c_s = 1.0$), where both KCIT and unconditional MMD fail, conditional MMD detected the direct effect at MMD $= 0.400$, and the integrated algorithm recovered the edge in all 10 seeds. This corresponds to a kernel estimator of Pearl's controlled direct effect (CDE).

This cascade of joint encoding (Phase 1), MMD complement (Phase 2), conditional MMD (Phase 3), and spurious edge removal (Phase 4) constitutes the design principle of the integrated algorithm (Algorithms 1–2). Each phase takes the output of its predecessor as input, progressively reinforcing and correcting the initial KCIT skeleton via MMD.

To assess how frequently path cancellation arises in practice, we measured the prevalence of cancellation candidates on the 12-variable mixed DAG of Benchmark 1E ($N = 500$, 10 seeds). A pair $(\mathrm{do}(X_i), X_j)$ is flagged as a cancellation candidate when the total MMD is near zero (< 0.02) but the conditional MMD, stratified by potential mediators, exceeds the total by at least a factor of 3. On average, 6.4 ± 3.1 cancellation candidates were detected per seed, and all 10 seeds contained at least one. A typical example is $\mathrm{do}(X_2) \to X_{10}$ with total MMD $= 0.004$ and conditional MMD $= 0.025$ (ratio $5.8\times$). This confirms that path cancellation is not a rare pathology but a systematic phenomenon in DAGs with ≥ 10 variables, validating the need for Phase 3 conditional MMD. Recent work on state-aware causal inference in physical systems (Martínez-Sánchez & Lozano-Durán, 2026) highlights the growing demand for methods that can handle nonlinear, state-dependent causal relationships.

6.2 Framework design

Pearl's do-operator is an idealized mathematical concept, but computational simulators widely used in physics and chemistry realize this idealization as $\mathcal{S}$-valid do-operators (Definition 1). Simulators guarantee intervention independence, mechanism preservation, and computational determinism in principle, free from the measurement error and selection bias inherent in real-world experiments. A bidirectional relationship is established: advances in computational science directly improve causal inference accuracy.

The KIIM ablation experiment (Table 7) revealed no significant difference between quantum-inspired kernel KIIM and classical RBF-KIIM, indicating that the structural advantage is concentrated at the KCIT (skeleton discovery) stage. The KCIT/KIIM separation architecture based on this finding is not merely an engineering compromise but has a structural rationale: the phase encoding of the quantum-inspired kernel requires sensitivity optimization specific to the HSIC test.

The applicability conditions for QCSI are: (i) the domain exhibits multiplicative or interaction-type dependencies, (ii) a simulator functioning as an $\mathcal{S}$-valid do-operator is available, and (iii) the number of variables is on the order of tens (qubit requirement $1 + L$). In systems dominated by linear SEMs, classical PC/GES remain superior and the return on investment of QCSI is low. The MMD integration (Phases 2–4) is effective on complex DAGs where KCIT's structural limitations are exposed (latent confounding, faithfulness violations), whereas for simple structures with $\sim$3 variables, KCIT + KIIM alone suffices.

6.3 Discovery accuracy governs estimation accuracy

The most important structural property of the QCSI pipeline is that discovery accuracy directly governs estimation accuracy. The distinction between confounders and mediators is determined by edge directions in the causal graph. Erroneously adjusting for mediators induces overadjustment bias.

In the Tox21 benchmark, DML adjusted for all covariates without directional information, causing overadjustment in 5 Mediated-type outcome variables and sign reversal in 4 of them. QCSI-Estimation, operating on directions correctly identified by QCSI-Discovery, estimated positive causal effects for all 12 outcome variables.

The estimation stage (Phase 4) is completed by Flow Matching and classical statistics, with no dependence on quantum-inspired kernels. The structural advantages of quantum-inspired kernels propagate indirectly to estimation via Discovery (Phases 1–3). This separation guarantees that changes in quantum-inspired kernel scaling do not affect estimation-stage stability.

A roadmap for NISQ hardware deployment and prospects for causal structure of quantum systems themselves (indefinite causal order, etc.) are given in Appendix J.

7 Conclusion

We proposed QCSI (Causal Structure Inference via Quantum-inspired kernels and Simulators), an integrated algorithm for causal discovery. QCSI is a framework that unifies three technical contributions.

The first contribution is the MMD-integrated algorithm (Phases 1–4). By progressively refining the initial KCIT skeleton via MMD on intervention data, it overcomes the structural limitations of KCIT (latent confounding, faithfulness violations). The direction determination component (KIIM) compares observational and interventional distributions via MMD and was correct for all pairs. Conditional MMD achieved edge recovery across all seeds under complete path cancellation, and spurious edge removal improved the quantum-inspired kernel's F1 by +0.022.

The second contribution is the verification of $\mathcal{S}$-valid do-operators. We stated three conditions under which a physics simulator functions equivalently to Pearl's do-operator and verified that three scientific simulators satisfy them. This verification justifies the intervention data used in Phases 2–4.

The third contribution is joint encoding. The ZZ phase $\phi = x_i \cdot x_j$ of the quantum-inspired kernel embeds variable products into quantum states, improving sensitivity to multiplicative dependencies in physical systems. On a multiplicative SEM, QCSI-ensemble outperformed IGSP by +0.243 in F1, and on Sachs synthetic data, it tied with UT-IGSP for the top score.

Through the integration of these contributions, QCSI correctly estimated positive causal effects for all 12 outcome variables in Tox21 real data, in settings where DML without directional information produced sign-reversed estimates. This confirms the structural property of the pipeline: improved discovery accuracy directly governs estimation accuracy.

8 Limitations

First, there is a tradeoff between skeleton F1 and direction accuracy. On synthetic laser physics data, QCSI-Discovery's skeleton F1 (0.527 ± 0.018) falls below classical PC (0.599 ± 0.019), while its direction accuracy

($64.8 \pm 4.6\%$) substantially exceeds PC's ($4.8 \pm 2.8\%$). Evaluation of causal discovery cannot rely on a single metric. systematic analysis of the skeleton–direction tradeoff is needed.

Second, the theoretical proof of quantum advantage has not been established at the computational complexity level. Proposition 2 is a structural argument about the representation power of quantum-inspired kernel MMD, which differs in nature from the rigorous proofs at the computational complexity level given by Liu et al. (2021) and Huang et al. (2022).

Third, all quantum-inspired kernel computations were performed on a classical simulator (Qiskit Aer StatevectorSampler), not on quantum hardware. Execution on NISQ devices and evaluation of quantum noise effects are beyond the scope of this paper.

Fourth, the MMD-integrated algorithm has limited applicability. The MMD integration (Phases 2–4) is effective on complex DAGs involving latent confounding and faithfulness violations, but for simple structures with $\sim$3 variables, KCIT + KIIM alone suffices. Additionally, in dense collider structures, Phase 3 may misidentify collider conditioning bias as a faithfulness violation, producing false edge recoveries. Improving Phase 3's collider detection conditions remains future work.

Fifth, comparison with standard benchmarks is limited. We compared with IGSP/UT-IGSP on synthetic data based on the Sachs et al. (Sachs et al., 2005) network structure, where QCSI-RBF tied with UT-IGSP at $\mathrm{F1} = 0.707$ (Table 5). However, evaluation on the actual Sachs flow cytometry data and on the DREAM challenge has not been conducted.

Sixth, ground truth causal graphs for real data are constrained. For HHG, SEI, and Tox21, we constructed ground truth DAGs based on domain knowledge (based on established physical and chemical theory (Corkum, 1993; Xu, 2004; Huang et al., 2016; Richard et al., 2016)) and evaluated skeleton F1 against them. However, these are partial graphs based on theoretical knowledge, not complete causal graphs.

Seventh, our $\mathcal{S}$-validity conditions assume that the simulator faithfully implements the causal mechanisms of the real system. When the simulator is misspecified (e.g., an approximate functional or missing physics), the intervention data may not correspond to the true do-operator, and the resulting causal graph may contain errors. Quantifying robustness to simulator misspecification is an important direction for future work.

Eighth, scalability to large DAGs has not been fully verified. Empirical timing on DAGs with $d = 10, 20, 30, 50$ variables showed that $d = 50$ requires 1,050 seconds per seed ($\sim$17 minutes). Extension to $d \geq 100$ will require conditioning set size restrictions or sparsity assumptions.

Broader Impact Statement

This work develops a methodological framework for causal discovery in scientific data using quantum-inspired kernels and physics simulators. The primary applications are in computational science domains (laser physics, battery chemistry, drug toxicity prediction) where causal understanding can improve material design and safety assessment. We do not foresee direct negative societal impacts from the methodological contributions of this paper. However, causal inference tools in general carry risks of misapplication when causal assumptions are violated or when results are over-interpreted as definitive causal claims from limited data. Users should exercise appropriate caution in applying QCSI to domains where causal conclusions have consequential policy implications.

References

C. Bannwarth, S. Ehlert, and S. Grimme. GFN2-xTB—an accurate and broadly parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics and density-dependent dispersion contributions. *Journal of Chemical Theory and Computation*, 15(3):1652–1671, 2019.

E. Bareinboim and J. Pearl. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences*, 113(27):7345–7352, 2016.

P. Brouillard, S. Lachapelle, A. Lacoste, S. Lacoste-Julien, and A. Drouin. Differentiable causal discovery from interventional data. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 2020.

Č. Brukner. Quantum causality. *Nature Physics*, 10:259–263, 2014.

M. Budišić, R. Mohr, and I. Mezić. Applied Koopmanism. *Chaos*, 22(4):047510, 2012.

Z. Chen, K. Zhang, L. Chan, and B. Schölkopf. Causal discovery via reproducing kernel Hilbert space embeddings. *Neural Computation*, 26(7):1484–1517, 2014.

V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21 (1):C1–C68, 2018.

P. B. Corkum. Plasma perspective on strong-field multiphoton ionization. *Physical Review Letters*, 71(13): 1994–1997, 1993.

K. Cranmer, J. Brehmer, and G. Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.

M. Deistler, P. J. Gonçalves, J. H. Macke, et al. Simulation-based inference: A practical guide, 2025. arXiv:2508.12939.

R. R. Dhanakshirur, Y. Wu, C. Guo, and K. Muandet. SCMD: A kernel-based distance for structural causal models to quantify transferability across environments, 2025. arXiv:2510.20481.

J. Fawkes and R. D. Shah. Doubly robust kernel statistics for testing distributional treatment effects. *Journal of Machine Learning Research*, 25:1–44, 2024.

J. Fuentes, A. Dehghani, and R. A. Brewster. Enhancing quantum machine learning: The power of non-linear optical reproducing kernels, 2024. arXiv:2407.13809.

T. Geffner, J. Antoran, A. Foster, W. Gong, C. Ma, E. Kiciman, A. Sharma, A. Lamb, M. Kukla, N. Pawlowski, M. Allamanis, and C. Zhang. Deep end-to-end causal inference, 2022. arXiv:2202.02195.

P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo, A. Dal Corso, S. de Gironcoli, S. Fabris, G. Fratesi, R. Gebauer, U. Gerstmann, C. Gougoussis, A. Kokalj, M. Lazzeri, L. Martin-Samos, N. Marzari, F. Mauri, R. Mazzarello, S. Paolini, A. Pasquarello, L. Paulatto, C. Sbraccia, S. Scandolo, G. Sclauzero, A. P. Seitsonen, A. Smogunov, P. Umari, and R. M. Wentzcovitch. QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials. *Journal of Physics: Condensed Matter*, 21(39):395502, 2009.

C. Giarmatzi and F. Costa. A quantum causal discovery algorithm. *npj Quantum Information*, 4:17, 2018.

C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969.

A. Gretton, K. Fukumizu, C. H. Teo, L. Song, B. Schölkopf, and A. Smola. A kernel statistical test of independence. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 20, pp. 585–592, 2007.

A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13:723–773, 2012.

Z. Guo and F. Dong. Linear causal discovery with interventional constraints. *Machine Learning*, 115:1–28, 2026.

V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567:209–212, 2019.

C. Heinze-Deml, J. Peters, and N. Meinshausen. Invariant causal prediction for nonlinear models. *Journal of Causal Inference*, 6, 2018.

H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill, and J. R. McClean. Quantum advantage in learning from experiments. *Science*, 376:1182–1186, 2022.

R. Huang, M. Xia, M.-H. Cho, S. Sakamuru, P. Shinn, K. A. Houck, D. J. Dix, R. S. Judson, K. L. Witt, R. J. Kavlock, R. R. Tice, and C. P. Austin. Tox21Challenge to build predictive models of nuclear receptor and stress response pathways. *Frontiers in Environmental Science*, 3:85, 2016.

A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W. Cross, B. R. Johnson, and J. M. Gambetta. Quantum computing with Qiskit, 2024. arXiv:2405.08810.

M. Kalisch and P. Bühlmann. Estimating high-dimensional directed acyclic graphs with the PC-algorithm. *Journal of Machine Learning Research*, 8:613–636, 2007.

H. Kawaguchi. Application of quantum computing to a linear non-Gaussian acyclic model for novel medical knowledge discovery. *PLOS ONE*, 18(4):e0283933, 2023.

T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.

B. O. Koopman. Hamiltonian systems and transformation in Hilbert space. *Proceedings of the National Academy of Sciences*, 17(5):315–318, 1931.

T. Kyono, Y. Zhang, and M. van der Schaar. CASTLE: Regularization via auxiliary causal graph discovery. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 2020.

M. Lewenstein, Ph. Balcou, M. Yu. Ivanov, A. L'Huillier, and P. B. Corkum. Theory of high-harmonic generation by low-frequency laser fields. *Physical Review A*, 49:2117–2132, 1994.

Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhatt, A. Stuart, and A. Anandkumar. Fourier Neural Operator for parametric partial differential equations. In *Proceedings of the Ninth International Conference on Learning Representations (ICLR)*, 2021.

Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le. Flow matching for generative modeling. In *Proceedings of the Eleventh International Conference on Learning Representations (ICLR)*, 2023. arXiv:2210.02747.

P. Lippe, T. Cohen, and E. Gavves. Efficient neural causal discovery without acyclicity constraints. In *Proceedings of the Tenth International Conference on Learning Representations (ICLR)*, 2022.

Y. Liu, S. Arunachalam, and K. Temme. A rigorous and robust quantum speed-up in supervised machine learning. *Nature Physics*, 17:1013–1017, 2021.

D. Lopez-Paz, K. Muandet, B. Schölkopf, and I. Tolstikhin. Towards a learning theory of cause-effect inference. In *Proceedings of the Thirty-Second International Conference on Machine Learning (ICML)*, pp. 1452–1461, 2015.

Y. Maeda, K. Arai, Y. Tanaka, Y. Terada, H. Ueno, and H. Tezuka. Quantum-enhanced causal discovery for a small number of samples. *Quantum Machine Intelligence*, 7:42, 2025. arXiv:2501.05007.

Á. Martínez-Sánchez and A. Lozano-Durán. Observational causality by states and interaction type for scientific discovery. *Communications Physics*, 9:15, 2026.

J. M. Mooij, S. Magliacane, and T. Claassen. Joint causal inference from multiple contexts. *Journal of Machine Learning Research*, 21(99):1–108, 2020.

K. Muandet, K. Fukumizu, B. Sriperumbudur, and B. Schölkopf. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends in Machine Learning*, 10(1–2):1–141, 2017.

K. Muandet, M. Kanagawa, S. Saengkyongam, and S. Marukatat. Counterfactual mean embeddings. In *Journal of Machine Learning Research*, volume 22, pp. 1–71, 2021.

J. Park, U. Shalit, B. Schölkopf, and K. Muandet. Conditional distributional treatment effect with kernel conditional mean embeddings and U-statistic regression. In *Proceedings of the Thirty-Eighth International Conference on Machine Learning (ICML)*, volume 139, pp. 8401–8412, 2021.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition, 2009.

J. Peters, P. Bühlmann, and N. Meinshausen. Causal inference using invariant prediction: Identification and confidence intervals. *Journal of the Royal Statistical Society: Series B*, 78(5):947–1012, 2016.

S. Rendle. Factorization machines. In *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, pp. 995–1000, 2010.

A. M. Richard, R. S. Judson, K. A. Houck, C. M. Grulke, P. Volarath, I. Thillainadarajah, C. Yang, J. Rathman, M. T. Martin, J. F. Wambaugh, T. B. Knudsen, J. Kancherla, K. Mansouri, G. Patlewicz, A. J. Williams, S. B. Little, K. M. Crofton, and R. S. Thomas. ToxCast chemical landscape: paving the road to 21st century toxicology. *Chemical Research in Toxicology*, 29:1225–1251, 2016.

K. Sachs, O. Perez, D. Pe'er, D. A. Lauffenburger, and G. P. Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308:523–529, 2005.

P. Sanchez and S. A. Tsaftaris. Diffusion causal models for counterfactual estimation. In *Proceedings of the First Conference on Causal Learning and Reasoning (CLeaR)*, 2022.

B. Schölkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, and Y. Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021.

M. Schuld. Supervised quantum machine learning models are kernel methods, 2021. arXiv:2101.11020.

R. Shaydulin and S. M. Wild. Importance of kernel bandwidth in quantum machine learning. *Physical Review A*, 106(4):042407, 2022.

X. Shen, F. Liu, H. Dong, Q. Lian, Z. Chen, and T. Zhang. Weakly supervised disentangled generative causal representation learning. *Journal of Machine Learning Research*, 23:1–55, 2022.

P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT Press, 2nd edition, 2000.

C. Squires, Y. Wang, and C. Uhler. Permutation-based causal structure learning with unknown intervention targets. In *Proceedings of the Thirty-Sixth Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 1039–1048, 2020.

N. Tancogne-Dejean, M. J. T. Oliveira, X. Andrade, H. Appel, C. H. Borca, G. Le Breton, F. Buchholz, A. Castro, S. Corni, A. A. Correa, U. De Giovannini, A. Delgado, F. G. Eich, J. Flick, G. Gil, A. Gomez, N. Helbig, H. Hübener, R. Jestädt, J. Jornet-Somoza, A. H. Larsen, I. V. Lebedeva, M. Lüders, M. A. L. Marques, S. T. Ohlmann, S. Pipolo, M. Rampp, C. A. Rozzi, D. A. Strubbe, S. A. Sato, C. Schäfer, I. Theophilou, A. Welden, and A. Rubio. Octopus, a computational framework for exploring light-driven phenomena and quantum dynamics of interacting particles. *Journal of Chemical Physics*, 152:124119, 2020.

S. Thanasilp, S. Wang, M. Cerezo, and Z. Holmes. Exponential concentration in quantum kernel methods. *Nature Communications*, 15:5200, 2024.

Y. Wang, L. Solus, K. Yang, and C. Uhler. Permutation-based causal inference algorithms with interventions. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, pp. 5822–5831, 2017.

M. O. Williams, I. G. Kevrekidis, and C. W. Rowley. A data-driven approximation of the Koopman operator: extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, 2015.

K. Xu. Nonaqueous liquid electrolytes for lithium-based rechargeable batteries. *Chemical Reviews*, 104: 4303–4417, 2004.

K. Xu. Electrolytes and interphases in Li-ion batteries and beyond. *Chemical Reviews*, 114(23):11503–11618, 2014.

H. Zenati, V. Mayer, A. Gretton, and K. Muandet. Doubly robust estimation of interventional mean embeddings, 2025. arXiv:2504.08840.

K. Zhang and A. Hyvärinen. On the identifiability of the post-nonlinear causal model. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 647–655, 2009.

K. Zhang, J. Peters, D. Janzing, and B. Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 804–813, 2012.

X. Zheng, B. Aragam, P. Ravikumar, and E. Xing. DAGs with NO TEARS: Continuous optimization for structure learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, 2018.

A Details of the Median Heuristic

The median heuristic does not change the content of the phase encoding. Rather, it optimizes the resolution at which the HSIC test reads the phase information. The ZZ phase $\phi = 2(\pi - x_i)(\pi - x_j)$ encodes the product $x_i \cdot x_j$ into the phase of the quantum state regardless of scaling. This is analogous to how Berry phase[7] $\gamma = \oint \langle\psi|\nabla_k|\psi\rangle dk$ captures topological invariants independent of parameterization. However, just as measuring Berry phase requires an integration loop of appropriate size relative to the Brillouin zone to avoid missing phase windings or aliasing, the HSIC test detecting phase structure through $k_Q = \cos^2(\Delta)$ requires proper calibration. The median heuristic automatically selects the operating point that maximizes statistical test sensitivity via the condition $\cos^2(\tilde{\Delta}) = 0.5$.

This principle shares a common foundation with the Kerr coherent state approach to continuous-variable quantum-inspired kernels (Fuentes et al., 2024). While the Kerr parameter controls curvature in feature space, the median heuristic achieves equivalent bandwidth control for discrete-qubit ZZFeatureMap with $O(N^2)$ preprocessing.

B Proofs of Theoretical Guarantees

B.1 Proof sketch of Proposition 2

The quantum-inspired kernel $k_Q(\mathbf{x}, \mathbf{x}') = |\langle\psi(\mathbf{x})|\psi(\mathbf{x}')\rangle|^2$ defines an embedding into a 2^n-dimensional Hilbert space. ZZFeatureMap encodes phases for subsets $|S| \leq 2$, explicitly representing variable products $x_i \cdot x_j$ in the feature space. In the kernel framework of Schuld (2021), this high-dimensional feature space reduces the mean squared error of MMD at rate $O(1/\sqrt{N})$ independently of the classical kernel. □

B.2 Proof sketch of Theorem 3

When $X_i \to X_j$ is true, the $\mathcal{S}$-valid do-operator $\mathrm{do}(X_i = k)$ changes the distribution of X_j (since $\beta \neq 0$). Conversely, $\mathrm{do}(X_j = k)$ does not change the distribution of X_i. Since MMD is a consistent estimator (Gretton et al., 2012), in the limit $N \to \infty$: $\mathrm{MMD}^2(P, Q) = 0 \iff P = Q$. □

B.3 Proof sketch of Theorem 4

By the concentration inequality of Gretton et al. (2012): $P(|\widehat{\mathrm{MMD}}^2 - \mathrm{MMD}^2| \geq \epsilon) \leq 2\exp(-m\epsilon^2/8\kappa^4)$. Setting $\epsilon = \Delta/2$ and applying to the error event yields Eq. (12). □

[7] M. V. Berry, "Quantal phase factors accompanying adiabatic changes," *Proc. Royal Soc. London A*, 392:45–57, 1984.

B.4 Proof sketch of Theorem 6

In the linear SEM $X = \alpha Z_{\text{conf}} + \eta_X$, $Y = \theta_{\text{causal}} X + \gamma Z_{\text{conf}} + \epsilon$, the Frisch–Waugh–Lovell theorem[8] implies that the difference between OLS and back-door adjusted estimates equals $\gamma \cdot \text{Cov}(Z_{\text{conf}}, X)/\text{Var}(X)$. Solving inversely yields Eq. (13). □

C Theoretical Analysis of Sample Size Dependence

We discuss why QCSI achieves optimal performance at $N = 200$ and why behavior changes at $N = 500$.

When the sample size N is large, linear correlations in multivariate data approach near-perfect levels ($r \to 1$), causing the numerical rank of kernel matrices used in KCIT to decline. In the residualized matrix $\tilde{K}_{X|\mathbf{Z}} = R_Z K_{X\mathbf{Z}} R_Z$, where $R_Z = \epsilon(K_Z + \epsilon I)^{-1}$, K_Z becoming approximately rank-1 degrades the numerical behavior of R_Z.

In practical terms,

- $N \leq 100$: insufficient statistical power
- $N = 200$: optimal balance
- $N \geq 500$: numerical degradation due to higher-order correlations

In real-world applications, $N \sim 100$–300 is typical, and QCSI demonstrates its advantage in this range.

D Details of Synthetic Benchmark 1A

Synthetic Benchmark 1A evaluates causal discovery accuracy on a 6-variable structural equation model (SEM) containing the multiplicative dependency $Y_1 = X_1 \cdot X_2^2$. The SEM is defined as follows:

$$X_1 \sim \mathcal{N}(0, 1) \tag{19}$$
$$X_2 \sim \mathcal{N}(0, 1) \tag{20}$$
$$X_3 \sim \mathcal{N}(0.1, 0.05) \tag{21}$$
$$Y_1 = X_1 \cdot X_2^2 + 0.1\epsilon_1 \tag{22}$$
$$Y_2 = Y_1 + X_3 + 0.1\epsilon_2 \tag{23}$$
$$Y_3 = f_{\text{nl}}(Y_2) + 0.1\epsilon_3 \tag{24}$$

where f_{nl} is a nonlinear activation function and $\epsilon_i \sim \mathcal{N}(0, 1)$ are independent noise terms. The key multiplicative structure is $Y_1 = X_1 \cdot X_2^2$: in an additive feature space (Gaussian kernel), the power to separate the independent contributions of X_1 and X_2 from their multiplicative interaction is limited. The true causal graph consists of 6 edges: $X_1 \to Y_1$, $X_2 \to Y_1$, $X_3 \to Y_2$, $Y_1 \to Y_2$, and $Y_2 \to Y_3$.

E Details of Synthetic Benchmark 1B

Synthetic Benchmark 1B tests whether intervention data can distinguish a v-structure ($T_1 \to M \leftarrow T_2$) from a chain ($T_1 \to M \to T_2$) in a 3-variable causal structure. For v-structures, the intermediate variable is generated as $M = f(T_1 \cdot T_2) + \epsilon$ with $\epsilon \sim \mathcal{N}(0, 0.3^2)$, where f is one of 20 different nonlinear multiplicative functions. For chains, the mechanism is additive: $M = 2T_1 + \epsilon$ and $T_2 = 2M + \epsilon$. Each setting uses $N = 200$ samples, 30 seeds, and significance level $\alpha = 0.02$.

Table 16 presents results for all 20 settings. The methods compared are the quantum-inspired kernel (Q), the ensemble of quantum-inspired and RBF kernels (E), RBF kernel alone (R), polynomial kernel $d = 2$ (P), and the classical PC and GES algorithms (using only observational data). Accuracy is measured as the percentage of seeds where the correct causal structure (v-structure or chain) is identified. Statistical significance is assessed by Wilcoxon signed-rank tests comparing Q against R across seeds.

[8] R. Frisch and F. V. Waugh, "Partial time regressions as compared with individual trends," *Econometrica*, 1(4):387–401, 1933.

Table 16: Accuracy (%) for all 20 settings of Benchmark 1B. Methods compared are quantum-inspired kernel (Q, ours), ensemble (E, ours), RBF (R, ours), polynomial $d = 2$ (P), classical PC, and GES. p-values are from Wilcoxon signed-rank tests (Q vs. R). $^{*}p < 0.05$, $^{**}p < 0.01$. Bold indicates the highest accuracy per row.

Regime	Causal relation $M = f(\cdot)$	Q	E	R	P	PC	GES	p(Q,R)
Quantum-inspired kernel significantly superior to RBF ($p < 0.05$)								
Phase trans.	$\tanh(5\,T_1T_2)$	**90.0**	**90.0**	75.0	50.0	50.0	50.0	0.020*
Wave	$\sin(\pi T_1)\sin(\pi T_2)$	**73.3**	66.7	48.3	48.3	50.0	50.0	0.0003**
Info-theoretic	$T_1T_2\log\lvert T_1T_2\rvert$	**66.7**	65.0	51.7	48.3	50.0	50.0	0.023*
Quantum-inspired kernel higher than RBF (not significant)								
Periodic	$3\sin(\pi\,T_1T_2)$	**96.7**	**96.7**	91.7	65.0	50.0	50.0	0.453
Coulomb	$T_1T_2/\sqrt{T_1^2+T_2^2}$	93.3	**95.0**	90.0	58.3	50.0	50.0	0.754
Heterosced.	$3\sin(\pi x)$+heterosced. ϵ	**98.3**	96.7	90.0	61.7	50.0	50.0	0.125
Regime switch	$\sin(\pi x)/x$ conditional	**96.7**	93.3	93.3	66.7	50.0	50.0	0.688
sin+exp	$2\sin(\pi x)+e^x-1$	**98.3**	96.7	96.7	68.3	50.0	50.0	1.000
poly+periodic	$2\,T_1T_2+2\sin(\pi x)$	**98.3**	96.7	96.7	75.0	50.0	50.0	1.000
sin·logistic	$\sin(\pi x)/(1+e^{-3x})$	**75.0**	**75.0**	70.0	51.7	50.0	50.0	0.581
Ensemble effective ($E \geq R$)								
Activation (elu)	$\mathrm{elu}(T_1T_2)$	90.0	**95.0**	**95.0**	71.7	50.0	50.0	0.453
Chaos	$4\,T_1T_2(1-T_1T_2)$	75.0	**88.3**	**88.3**	63.3	58.3	53.3	0.046
GARCH	$T_1T_2(1+2\lvert T_1T_2\rvert)$	80.0	**96.7**	**96.7**	93.3	50.0	50.0	0.006
Ensemble as safety net ($Q < R$ but $E \approx R$)								
Damped osc.	$e^{-r^2/2}\cos(\pi\,T_1T_2)$	81.7	**96.7**	**96.7**	96.7	50.0	50.0	0.023
Morse	$(1-e^{-2\,T_1T_2})^2$	68.3	**95.0**	**95.0**	70.0	55.0	56.7	0.0003
Bloch sphere	Bloch correlation	56.7	**86.7**	**86.7**	78.3	51.7	51.7	0.0001
Hill function	$\lvert x\rvert/(0.5+\lvert x\rvert)$	60.0	**95.0**	**95.0**	95.0	50.0	50.0	0.000
Gabor	$e^{-r^2/2}\cos(5\,T_1T_2)$	63.3	**96.7**	**96.7**	96.7	50.0	50.0	0.000
Higher-order	$\sin(\pi\,T_1^2T_2)$	75.0	**96.7**	**96.7**	96.7	51.7	53.3	0.002
Step	$\mathrm{sign}(T_1T_2)$	48.3	48.3	48.3	48.3	50.0	50.0	1.000

The results in Table 16 reveal a clear dependence of quantum-inspired kernel performance on the nonlinear structure of the causal relationship. On phase-transition, wave, and information-theoretic structures, the quantum-inspired kernel is significantly superior to RBF ($p < 0.05$). These structures involve sharp thresholds (tanh), oscillatory interference ($\sin\cdot\sin$), or logarithmic scaling ($x\log|x|$), where the $\cos^2(\Delta)$ kernel's periodic sensitivity provides a detection advantage. On periodic multiplicative, Coulomb, and regime-switching structures, the quantum-inspired kernel outperforms RBF but without reaching statistical significance ($p > 0.1$), indicating that the advantage is present but modest. On damped oscillation, Morse, Hill function, and Gabor structures, the quantum-inspired kernel alone underperforms RBF. These structures involve amplitude modulation or saturation that attenuates the multiplicative signal. In these cases, the ensemble strategy maintains RBF-level performance, acting as a safety net that prevents the quantum-inspired kernel's lower accuracy from degrading overall results. Across all 20 settings, intervention-based methods (Q, E, and R) are significantly superior to the purely observational methods PC and GES (both at approximately 50% accuracy) at $p < 0.01$, confirming that intervention data is essential for distinguishing v-structures from chains.

The 20 regimes of Benchmark 1B were designed to simulate various nonlinear multiplicative structures that arise in physical and chemical systems, and we now describe the regime correspondence with real data. Each regime corresponds to a specific class of nonlinear mechanisms encountered in the real-data benchmarks (Benchmarks 2–4). Establishing this correspondence serves two purposes: it validates that the synthetic benchmark covers the functional forms relevant to real applications, and it helps interpret which regimes predict performance on specific real-data domains.

In the context of Benchmark 2 (HHG laser physics), Table 17 lists the regime correspondences. The periodic regime directly models the multiplicative product structure $U_p \propto I \cdot \lambda^2$ central to the ponderomotive energy. The wave regime captures the interference between short and long quantum trajectories that determines harmonic phase. The Coulomb regime represents the atomic Coulomb potential governing tunnel ionization rates. The phase-transition regime models the steep sigmoid transition from tunneling to barrier-suppression ionization as laser intensity increases. The poly+periodic regime captures the combined effect of cutoff energy (polynomial in U_p) and spectral structure (periodic in harmonic order).

Table 17: Regime correspondence between Benchmark 1B and Benchmark 2 (HHG laser physics).

Priority	Regime	Rationale
Primary	Periodic	$U_p \propto I \cdot \lambda^2$ is a direct product structure; HHG spectra are periodic in odd harmonics
Primary	Wave	Harmonic phase is determined by wave interference of quantum trajectories (short and long)
Primary	Coulomb	Atomic Coulomb potential governing tunnel ionization
Secondary	Phase trans.	Tunnel $\rightarrow$ barrier-suppression ionization transition is a steep sigmoid
Secondary	poly+periodic	Polynomial + periodic mixture corresponds to cutoff + spectral structure

For Benchmark 3 (SEI battery chemistry), Table 18 lists the regime correspondences. The phase-transition regime models the threshold-like SEI formation process, where a critical reduction potential must be exceeded before stable passivation begins. The activation (elu) regime represents the Arrhenius-type reaction rate dependence, which shows exponential increase with electrochemical driving force followed by saturation as the SEI layer thickens. The Coulomb regime captures the electrostatic contribution to the electrochemical reduction potential. The Morse regime models bond dissociation of electrolyte molecules, which follows a Morse-type potential in the initial decomposition step. The regime-switch regime reflects that different additive species (carbonates, sulfones, phosphate esters) undergo qualitatively different decomposition pathways.

Table 18: Regime correspondence between Benchmark 1B and Benchmark 3 (SEI battery chemistry).

Priority	Regime	Rationale
Primary	Phase trans.	SEI formation is threshold-like: a sigmoid transition once the critical reduction potential is exceeded
Primary	Activation (elu)	Arrhenius-type reaction rate: exponential increase with driving force, saturating
Secondary	Coulomb	Electrostatic driving force of electrochemical reduction
Secondary	Morse	Bond dissociation of electrolyte molecules follows a Morse-type potential
Secondary	Regime switch	Reaction regime differs by additive species (carbonate/sulfone/phosphate ester)

For Benchmark 4 (Tox21 molecular toxicity), Table 19 lists the regime correspondences. The Hill function regime directly models the dose-response relationship in receptor binding, where the activity follows $|x|/(K + |x|)$ and the Hill coefficient determines the steepness of the sigmoidal response. The regime-switch regime captures the qualitative difference between the 7 Direct endpoints (nuclear receptor pathways) and

the 5 Mediated endpoints (stress response pathways), which exhibit different functional forms and different causal pathway structures. The phase-transition regime models the sigmoidal receptor activation and inactivation curves that characterize nuclear receptor signaling. The info-theoretic regime represents the entropic information processing in biological signal transduction cascades. The higher-order regime captures the interaction terms between the structural descriptor T (aromatic ring count) and electronic descriptors Z_i (HOMO, LUMO) that jointly determine biological activity.

Table 19: Regime correspondence between Benchmark 1B and Benchmark 4 (Tox21 molecular toxicity).

Priority	Regime	Rationale
Primary	Hill function	Receptor binding dose-response follows a Hill function $\|x\|/(K+\|x\|)$
Primary	Regime switch	12 endpoints split into Direct (7) and Mediated (5) with different functional forms
Secondary	Phase trans.	Receptor activation/inactivation is sigmoidal
Secondary	Info-theoretic	Entropic information processing in biological signal transduction
Secondary	Higher-order	Higher-order interactions between structural descriptor T and electronic descriptors Z_i

F Details of Synthetic Benchmark 1E

We illustrate the step-by-step operation of the integrated algorithm (Phases 1–4) on a 12-variable mixed DAG (seed=0, $N = 500$, RBF kernel). The true DAG has 16 edges, and latent confounders U_1, U_2, U_3 induce spurious dependencies among observed variables.

F.1 Phase 1: KCIT skeleton

The PC algorithm with KCIT constructs the initial skeleton (Table 20).

Table 20: Phase 1 KCIT skeleton (seed=0, $N = 500$, RBF kernel)

Category	Edges
True positives (12/16)	$(0,1),(0,2),(0,4),(1,3),(1,4),(2,3),(3,5),(4,7),(5,8),(6,9),(7,9),(8,10)$
False positives (4)	$(3,4),(4,9),(5,10),(2,11)$
Missing (4/16)	$(4,6),(6,7),(9,11),(10,11)$

Skeleton F1 $= 2 \times 12/(16+16) = 0.750$. Among the missing edges, $(4,6)$ and $(6,7)$ are caused by latent confounder U_2, and $(9,11)$ and $(10,11)$ by latent confounders U_3 and U_4; these are structural limitations that persist regardless of N.

F.2 Phase 2: MMD intervention effect map

The effect of do(X_i=0) ($i \in \{1,2,3,4,5\}$) is measured via MMD across all variables. Only significant ($p < 0.05$) MMD values are shown (— denotes the intervention target itself; · denotes non-significant).

Two key observations emerge from Tables 21 and 22.

Quantification of causal distance: The effect of do(X_1) is largest for direct children (X_2: 0.650, X_5: 0.629) and attenuates for indirect descendants (X_8: 0.456, X_7: 0.271). MMD values reflect causal distance in the DAG.

Table 21: MMD intervention effect map (X_1–X_6). seed=0, $N = 500$, RBF kernel.

	X_1	X_2	X_3	X_4	X_5	X_6
do(X_1)	—	.650	.026	.204	.629	.031
do(X_2)	·	—	·	.361	.245	.051
do(X_3)	·	·	—	.060	·	·
do(X_4)	·	·	·	—	·	.087
do(X_5)	·	·	·	·	—	·

Table 22: MMD intervention effect map (X_7–X_{12}). Continuation of Table 21.

	X_7	X_8	X_9	X_{10}	X_{11}	X_{12}
do(X_1)	.271	.456	.023	.263	·	.041
do(X_2)	.107	.199	.033	.102	·	.026
do(X_3)	·	·	·	·	·	·
do(X_4)	·	·	.054	·	.024	·
do(X_5)	.297	.474	·	.287	·	.021

Direction identification: do(X_1) $\rightarrow X_2$ has MMD = 0.650 while do(X_2) $\rightarrow X_1$ is non-significant ($\cdot$), giving a ratio $> 300\times$ that clearly identifies $X_1 \rightarrow X_2$ as the causal direction. Direction identification was correct for all 9 pairs (100% accuracy).

F.3 Phase 3: Conditional MMD

We examine the missing edge $(0, 4)$ (i.e., X_1–X_5). do(X_1) $\rightarrow X_5$ has MMD = 0.629 > 0.01, which does not satisfy Phase 3's firing condition (total effect ≈ 0). This is a missing edge due to latent confounding, not path cancellation, so Phase 3 correctly does not fire (for an example of faithfulness violation recovery, see Table 12 in the main text).

F.4 Phase 4: Spurious edge removal

Both directions of intervention effect are examined for the 4 false positive edges (Table 23).

Table 23: Phase 4 spurious edge removal decisions

False edge	do(i)$\rightarrow j$	do(j)$\rightarrow i$	Decision	Reason
$(3, 4)$	MMD= 0.002, p=0.35	MMD= 0.002, p=0.30	Remove	Both non-significant
$(4, 9)$	Int. data available	*No int. data*	Cannot verify	No intervention on X_9
$(5, 10)$	*No int. data*	*No int. data*	Cannot verify	No intervention on either
$(2, 11)$	Int. data available	*No int. data*	Cannot verify	No intervention on X_{11}

Edge $(3, 4)$ has non-significant MMD in both directions ($p > 0.05$), indicating no causal path. This edge reflects a spurious dependency induced by U_1, which intervention data reveals as non-causal. The remaining 3 edges cannot be verified due to the constraint that intervention targets are limited to $\{X_1, ..., X_5\}$.

F.5 Integrated result

As shown in Table 24, Phase 4 removed one false edge, improving F1 by +0.024. Direction identification was correct for all detected edges (100%). Expanding intervention targets to X_6–X_{12} would enable verification of the remaining 3 false edges.

Table 24: Integrated algorithm result (Phase 1 vs. Phases 1–4)

	Edge count	F1
Phase 1 (KCIT skeleton)	16 (12 TP + 4 FP)	0.750
Phases 1–4 (MMD integrated)	15 (12 TP + 3 FP)	**0.774**

G Details of Benchmark 2

Benchmark 2 tests v-structure identification using high-harmonic generation (HHG) simulation data. This section describes the physical background, variable definitions, computational setup, and physical interpretation, providing sufficient detail for full reproducibility of the HHG experiments reported in §5.2.1.

High-harmonic generation is an ultrafast nonlinear optical process in which odd-order harmonics of the incident laser frequency are emitted from atoms or molecules in a strong laser field. The underlying mechanism is described by the semiclassical three-step model (Corkum, 1993): (1) the valence electron undergoes tunnel ionization through the potential barrier suppressed by the laser field, (2) the freed electron is accelerated by the oscillating field and gains kinetic energy, and (3) the electron recollides with the parent ion upon field reversal, releasing the accumulated kinetic energy as a high-energy photon. The maximum photon energy (cutoff) is given by

$$E_{\text{cutoff}} = I_p + 3.17\,U_p \tag{25}$$

where I_p is the ionization potential and U_p is the ponderomotive energy, defined as

$$U_p\,[\text{eV}] = 9.33 \times \frac{I}{10^{14}} \times \left(\frac{\lambda}{1000}\right)^2 \tag{26}$$

with I the laser intensity in W/cm^2 and λ the wavelength in nm. The product $I \cdot \lambda^2$ constitutes the multiplicative dependency central to this benchmark.

The mapping between abstract variables in the main text (§5.2.1) and domain-specific physical quantities is as follows. T_1 (treatment 1) corresponds to the internuclear distance R in atomic units, which serves as the intervention target via do(R) (fixed at specified values in the TDDFT calculation). T_2 (treatment 2) corresponds to the laser field amplitude E_0 in atomic units, where $E_0^2 \propto I$. M (intermediate variable) corresponds to the ponderomotive energy U_p in eV, which exhibits the multiplicative structure $U_p \propto I \cdot \lambda^2$. Z (confounding candidate) corresponds to the ionization potential I_p in eV. Y_1 (outcome 1) corresponds to the harmonic intensity H_{int} in arbitrary units, and Y_2 (outcome 2) corresponds to the harmonic phase H_{phase} in radians. In observational data, T_1 and T_2 are negatively correlated, making the v-structure $T_1 \to M \leftarrow T_2$ and the chain $T_1 \to M \to T_2$ indistinguishable from conditional independence tests alone. Intervention via TDDFT simulation (do(T_1)) breaks this symmetry and enables structural identification.

Time-dependent density functional theory (TDDFT) calculations were performed using the Octopus package (Tancogne-Dejean et al., 2020), a real-space, real-time TDDFT code. The computational parameters were chosen to balance accuracy and throughput: real-space grid spacing of 0.4 a.u., spherical computational domain with radius 30 a.u., time step of 0.005 a.u., total propagation time of $T = 800$ a.u., and the local density approximation (LDA) for the exchange-correlation functional. The SIC-ADSIC self-interaction correction was applied to improve ionization potential accuracy. Sample sizes were 200 for observational data and 70 for cross-condition interventional data. Cross-condition interventions generate (E_0, R) combinations that lie outside the range of the observational dataset, ensuring that the intervention data provide genuinely new information about the causal structure.

The v-structure $E_0 \to U_p \leftarrow R$ detected by QCSI-Discovery is consistent with the three-step model prediction: both the field amplitude E_0 (via $I \propto E_0^2$) and the wavelength parameter (controlled through R) independently determine U_p. The conditional independence $U_p \perp\!\!\!\perp I_p \mid \{E_0, R\}$ reflects the physical fact that the ionization potential I_p contributes additively to E_{cutoff} but does not affect U_p, which depends only on the laser parameters. QCSI-Discovery correctly recovered the causal sign of the $E_0 \to H_{\text{int}}$ effect (positive: +5.921 eV/a.u., $R^2 = 0.983$), confirming that stronger laser fields produce higher harmonic yields, as predicted by the three-step model.

H Details of Benchmark 3

Benchmark 3 tests causal effect estimation using electronic structure calculation data for lithium-ion battery solid electrolyte interphase (SEI) additives. This section describes the electrochemical background, variable definitions, computational setup, confounding analysis, and design implications, providing sufficient detail for full reproducibility of the SEI experiments reported in §5.2.2.

The solid electrolyte interphase (SEI) in lithium-ion batteries forms on the anode surface through reductive decomposition of electrolyte additives during the initial charging cycles (Xu, 2014). The thermodynamic driving force for this decomposition is determined by the additive's LUMO (lowest unoccupied molecular orbital) energy level: a lower LUMO implies a higher electron affinity and thus easier electrochemical reduction. The resulting SEI film passivates the anode surface, preventing further electrolyte decomposition and controlling the long-term capacity retention of the battery. Understanding the causal relationship between additive electronic structure and battery performance is essential for rational electrolyte design.

The mapping between abstract variables in the main text (§5.2.2) and domain-specific quantities is as follows. The treatment T is the LUMO level in eV, which serves as an indicator of reduction susceptibility. The confounders are Z_1 (HOMO level in eV, an indicator of oxidative stability) and Z_2 (band gap in eV, approximately $Z_2 \approx Z_1 - T$). The outcome Y is the capacity retention in percent, measuring long-term battery performance after repeated charge-discharge cycling. The causal mechanism $T \to Y$ represents the pathway in which a low LUMO leads to easy reduction, which in turn promotes stable SEI formation and hence high capacity retention. This mechanism manifests as a negative correlation $r(T, Y) = -0.853$ in the observational data.

The additive dataset consists of 20 entries: 6 additives with DFT-computed electronic properties and 14 solvents and additional additives with literature values for electronic descriptors and experimentally measured capacity retention. Table 25 lists the 6 DFT-computed additives with their HOMO levels, LUMO levels, and literature capacity retention values. The additives are: EC (ethylene carbonate), FEC (fluoroethylene carbonate), VC (vinylene carbonate), LiBOB (lithium bis(oxalate)borate), DFEC (difluoroethylene carbonate), and TEP (triethyl phosphate). Literature sources for capacity retention values are Xu (2014) for EC, FEC, VC, and LiBOB; Xue et al. (2021, ACS Applied Materials & Interfaces) for DFEC; and Kim et al. (2025, Advanced Energy Materials) for TEP.

Table 25: DFT-computed properties and literature capacity retention for SEI additives

Additive	HOMO (eV)	LUMO (eV)	Cap. retention (%)
EC	−7.829	−0.866	78
FEC	−8.124	−0.987	92
VC	−7.412	−0.589	88
LiBOB	−6.945	−1.234	95
DFEC	−7.628	−1.003	85
TEP	−5.929	−0.679	75.5

The electronic structure calculations were performed using the Quantum ESPRESSO package (Giannozzi et al., 2009), a plane-wave pseudopotential DFT code. The computational parameters were: the Perdew–Burke–Ernzerhof (PBE) generalized gradient approximation for the exchange-correlation functional, a plane-wave kinetic energy cutoff of 60 Ry (with a charge density cutoff of 600 Ry), a vacuum layer of 15 Å to prevent periodic image interactions, Gamma-point-only k-point sampling (appropriate for isolated molecular calculations), and PAW (projector augmented wave) pseudopotentials. The simulator computes LUMO levels independently for each molecular structure, satisfying all three $\mathcal{S}$-validity conditions: the input (molecular geometry) fully specifies the output (electronic properties) without dependence on external experimental conditions.

Under a linear structural equation model, electronic properties other than LUMO (specifically HOMO, dipole moment, and polarizability) act as confounders of the $T \to Y$ relationship. The $\hat{\gamma}$ estimator (Theorem 6, Eq. 5) yields $\hat{\gamma} = 0.49$, which exceeds the crossover point (~ 0.25) identified in Benchmark 1C. This places

the SEI dataset in the regime where observational methods alone cannot recover the true causal effect, and interventional correction via QCSI-Estimation is necessary.

The causal effect of $-32.6\%/\text{eV}$ estimated by QCSI-Estimation has concrete implications for battery additive design. It suggests that lowering the LUMO level by 0.3–0.5 eV can improve capacity retention by approximately 10–16%. Specific chemical strategies include fluorination (the EC $\rightarrow$ FEC modification lowers LUMO by -0.121 eV, and the FEC $\rightarrow$ DFEC modification by a further -0.016 eV) and the introduction of electron-withdrawing groups. These guidelines connect the causal analysis to actionable molecular design decisions.

I Details of Benchmark 4

Benchmark 4 tests causal effect estimation and method crossover analysis using the Tox21 toxicity prediction dataset (Huang et al., 2016), which contains 6,097 molecules tested across 12 toxicity endpoints. This section describes the toxicological background, variable definitions, computational setup, evaluation metrics, and implications for drug design, providing sufficient detail for full reproducibility of the Tox21 experiments reported in §5.2.3.

The Tox21 dataset (Huang et al., 2016) is a publicly available benchmark from the Toxicology in the 21st Century initiative. It tests each molecule against 12 biological endpoints, grouped into two categories: seven nuclear receptor pathways (NR-AR, NR-AR-LBD, NR-AhR, NR-Aromatase, NR-ER, NR-ER-LBD, NR-PPAR-gamma) and five stress response pathways (SR-ARE, SR-ATAD5, SR-HSE, SR-MMP, SR-p53). Each endpoint is a binary activity label indicating whether the molecule activates the corresponding biological pathway. The dataset provides a realistic setting for evaluating causal discovery and effect estimation in high-dimensional molecular data where multiple causal pathways coexist.

The mapping between abstract variables in the main text (§5.2.3) and domain-specific quantities is as follows. The treatment T is the number of aromatic rings (AromaticRings), a fundamental structural feature of organic molecules that influences both electronic properties and biological activity. The confounding candidates $Z_1, \ldots, Z_k$ are electronic and physicochemical descriptors: HOMO and LUMO levels (from xTB calculation), LogP (lipophilicity), molecular weight (MolWt), and topological polar surface area (TPSA). The outcomes $Y_1, \ldots, Y_{12}$ are the 12 toxicity endpoint activity labels. Mediators M are a subset of the confounding candidates that form $T \rightarrow M \rightarrow Y_j$ paths, meaning they are causally downstream of T. The propagation of $\text{do}(T)$ effects to all descriptor variables ($p < 0.001$ for each) reflects the physical reality that changing aromatic ring count simultaneously alters the molecular electronic structure (HOMO and LUMO levels) and physicochemical properties (LogP and MolWt). The distinction between Direct effects ($T \rightarrow Y_j$) and Mediated effects ($T \rightarrow M \rightarrow Y_j$) identified by QCSI-Discovery is consistent with established structure-activity relationships in computational toxicology.

Electronic structure calculations were performed using GFN2-xTB (Bannwarth et al., 2019), a semi-empirical tight-binding method that provides accurate electronic descriptors at a fraction of the cost of full DFT calculations. Molecular geometries were optimized using xTB default convergence criteria. The computed electronic descriptors include HOMO and LUMO energies, HOMO-LUMO gap, dipole moment, ionization potential, and electron affinity. Structural descriptors (molecular weight, aromatic ring count, heteroatom count, rotatable bond count, and polar surface area) were computed using RDKit. GFN2-xTB functions as an $\mathcal{S}$-valid do-operator: it deterministically computes electronic descriptors (output) from molecular structure (input), with the calculation for each molecule being independent of all other molecules and external experimental conditions, satisfying condition 3 (computational determinism).

For evaluation, the causal effect θ_e from electronic descriptor X to toxicity endpoint Y_e is estimated for each endpoint $e \in \{1, \ldots, 12\}$. Bias is defined as the relative error (in percent) with respect to the reference estimate produced by QCSI-Estimation, which uses the correct adjustment set derived from the QCSI-Discovery causal graph. Table 15 in the main text reports the 12-endpoint average bias for each estimation method. DML (Double Machine Learning (Chernozhukov et al., 2018)), which adjusts for all covariates without directional information, erroneously adjusts for mediators in 5 of 12 endpoints, producing sign-reversed estimates in 4 of them. This overadjustment bias is a direct consequence of the lack of causal

direction information, confirming the structural property of the QCSI pipeline: discovery accuracy governs estimation accuracy.

The results have implications for drug design strategy. In the quantitative structure-activity relationship (QSAR) literature, combinations of electronic descriptors (HOMO and LUMO levels) and structural descriptors (aromatic rings and others) are considered important predictors of toxicity. Our causal analysis adds directionality to these associations: the LUMO level emerges as the primary causal descriptor for nuclear receptor binding, while the HOMO level merely exhibits correlation due to collinearity with LUMO. This implies that molecular modifications aimed at reducing receptor binding should target LUMO rather than HOMO in drug design and lead optimization.

J Future Outlook

J.1 Roadmap for NISQ hardware deployment

Current NISQ devices already provide the tens-of-qubit scale required for QCSI deployment (IBM Heron r3: 156 qubits; IonQ Forte: 36 qubits; QuEra Aquila: 256 qubits). QCSI's qubit requirement is $1 + L$ (L: maximum conditioning set size); for the HHG experiment ($d = 6$), 3 qubits suffice. The main challenge for hardware deployment is degradation of kernel matrix accuracy due to quantum noise, requiring combination with error mitigation (e.g., zero-noise extrapolation).

A short-to-medium-term roadmap envisions: 1–2 years for execution on IBM Quantum with noise sensitivity assessment; 3–5 years for QCSI with error mitigation on scientific data; 5–10 years for quantum advantage demonstration on error-corrected devices.

J.2 Extension to causal structure of quantum systems

This work applies quantum-inspired kernels to classical scientific data, but extension to the causal structure of quantum systems themselves is also within scope. Brukner's (Brukner, 2014) theory of indefinite causal order and the quantum process matrix formalism of Giarmatzi & Costa (2018) suggest that causal order itself may exist in quantum superposition. In physical processes under strong laser fields such as HHG, multiple quantum trajectories of the electron interfere to determine observables. The correspondence between causal structure extracted by QCSI from classical data and the quantum causal structure described by process matrices remains an unexplored research area.

J.3 Koopman operators and nonlinear causal inference

The Koopman operator (Koopman, 1931; Budišić et al., 2012) $\mathcal{K}$ acts on observables g of a dynamical system $x_{t+1} = F(x_t)$ as $\mathcal{K}g(x) = g(F(x))$, lifting nonlinear dynamics into a linear (but infinite-dimensional) operator on a function space. Finite-dimensional approximations such as Extended Dynamic Mode Decomposition (Williams et al., 2015) are equivalent to RKHS embeddings, providing a direct mathematical connection to kernel-based causal inference.

The ZZFeatureMap used in our joint encoding maps variable products $x_i \cdot x_j$ into quantum state phases, which can be interpreted as a finite-dimensional lifting analogous to Koopman eigenfunctions. Both approaches share the principle of linearizing nonlinear structure through a change of representation space.

For time-series causal discovery (e.g., battery degradation across charge-discharge cycles), Koopman-space representations could enable Granger causality (Granger, 1969) or VAR-based methods to detect nonlinear causal relationships that are invisible in the original state space. The current QCSI framework operates on static (cross-sectional) data, and extending it to Koopman-based temporal causal inference is a direction for future work.