SHADOWFLUID: TASK-ALIGNED REDUCED DYNAMICS FOR QUANTUM FLUID SIMULATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Quantum fluid simulation typically evolves the full quantum state, even when downstream tasks require only a small set of coarse-grained observables. We present SHADOWFLUID, a task-aligned, operator-first framework that constructs and evolves a reduced representation sufficient for target observables such as low-frequency density and energy statistics. Given a Hamiltonian and target observables, SHADOWFLUID builds a multi-reference dictionary of Fourier-mode coherence operators whose expectations form a reduced density-matrix block governed by a low-dimensional shadow generator. The reduced dynamics are exact when the dictionary is invariant under commutation with the Hamiltonian; otherwise, we derive a truncated evolution with a state-independent commutator-leakage indicator computable before simulation. We further establish a Frobenius-norm error bound with linear-in-time growth. Experiments on two-dimensional Schrödinger-flow benchmarks show exact recovery of the low-pass dynamics in the closed regime and sub-1% density error under nontrivial mode coupling at moderate cutoffs, with substantially smaller errors on target observables. These results demonstrate that task-relevant fluid dynamics can be extracted through compact, error-controlled representations without evolving and reconstructing full quantum state.

1 INTRODUCTION

Simulating fluid dynamics at high resolution remains computationally intensive due to the wide range of spatial and temporal scales involved. Quantum computing offers a potential path toward accelerating parts of this workload Lloyd (1996); Harrow et al. (2009); Georgescu et al. (2014); Babbush et al. (2023), and recent NISQ-era Preskill (2018) demonstrations Meng et al. (2024) have realized end-to-end digital simulation of 2D flows on a superconducting processor via a hydrodynamic Schrödinger formulation. Significant obstacles persist, however: full-field characterization can require exponentially many measurement shots, raising the question of whether useful flow statistics can be extracted without full-state reconstruction Aaronson (2018); Huang et al. (2020). In quantum CFD, the quantum state encodes an entire discretized flow field, yet many downstream objectives are far more compact. Practitioners typically require low-frequency density structure, energy content, or other coarse statistics rather than a full-field reconstruction at every time step. This mismatch motivates an operator-first strategy Mezić (2005): evolve only a small, task-selected set of operators instead of the entire wavefunction. Our approach builds on shadow Hamiltonian simulation Somma et al. (2025). Given a Hamiltonian and an operator set, shadow Hamiltonian simulation investigates conditions under which the operator expectations can be obtained from initial expectations alone. A sufficient condition is an *invariance property* (IP): commutators of the Hamiltonian with each operator remain within the span of the operator set, inducing a reduced-dimensional *shadow Hamiltonian* (§ 2.4). When this shadow Hamiltonian is Hermitian, the shadow state evolves in the reduced dimension rather than the full Hilbert space.

We present ShadowFluid, an operator-first framework for quantum fluid simulation. Starting from a set of task operators (e.g.,low-frequency density or energy observables), ShadowFluid constructs a multi-reference coherence dictionary in the Fourier basis. The dictionary tracks a sub-block of the density matrix that captures all task-relevant expectations. For diagonal Hamiltonians (zero potential), the invariance property holds exactly and the shadow evolution reproduces the full low-pass baseline to machine precision. For nontrivial potentials, the dictionary yields a truncated approxi-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

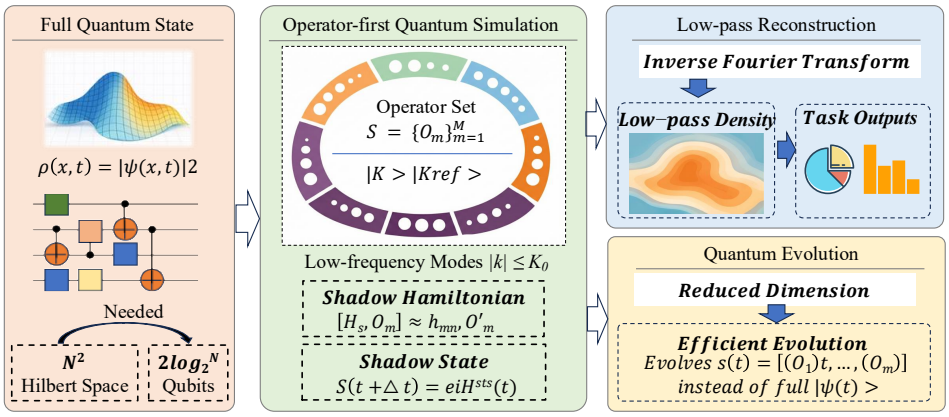


Figure 1: Overview of the ShadowFluid pipeline. From a full quantum state in high-dimensional Hilbert space, ShadowFluid selects a task-aligned operator dictionary, evolves the expectation matrix in reduced dimension, and reconstructs low-pass task outputs.

mation Quarteroni et al. (2015) whose accuracy is governed by an a priori computable commutator leakage residual. Our contributions are as follows.

- We formalize task-driven quantum fluid simulation through a coherence dictionary (§ 3.1), where task-relevant modes are collected and extended by coupling-graph closure on the Hamiltonian's Fourier-space connectivity.
- We establish a Frobenius-norm error bound that grows linearly in time because unitary conjugation preserves the Frobenius norm (§ 3.3). The leakage residual serves as a computable quality indicator that can be evaluated before any simulation.
- We evaluate ShadowFluid on 2D Schrödinger-flow benchmarks with cosine potentials across a wide range of coupling strengths (§ 4). The experiments confirm that the a priori bound holds uniformly, error grows linearly in time, and density error is dominated by low-pass truncation rather than the dictionary approximation.

2 Background and Preliminaries

2.1 Representation of Fluid Dynamics

Wavefunction and Velocity. We represent a 2D flow state by a complex wavefunction $\psi(x,t) \in \mathbb{C}$ defined over a spatial domain $x = (x_1, x_2) \in \Omega$. The associated density is $\rho(x,t) = |\psi(x,t)|^2$. Writing $\psi(x,t) = \sqrt{\rho(x,t)}\, e^{i\theta(x,t)}$ (Madelung transform Madelung (1927)), the phase θ induces a velocity field (up to problem-dependent constants) $v(x,t) \propto \nabla\theta(x,t)$, so that ρ and θ jointly encode fluid-like state variables.

Schrödinger-Type Dynamics. The evolution follows a Schrödinger-type PDE $i\,\partial_t\psi(x,t) = \Big(-\frac{1}{2}\nabla^2 + V(x,t)\Big)\psi(x,t)$, where the kinetic term governs dispersion in Fourier space and V captures problem-specific effects (external forcing, effective potentials, or multi-component couplings). The free case $V \equiv 0$ serves as the baseline and a building block for operator splitting.

2.2 Quantum Fourier Transform (QFT)

Grid Discretization and State Encoding. Let the spatial grid have $N = 2^n$ points per dimension (total N^2 points). We index grid points by integers $x = (x_1, x_2)$ with $x_d \in \{0, \ldots, N-1\}$ and encode $\psi(x,t)$ as amplitudes of a computational basis state $|\psi(t)\rangle = \sum_{x_1=0}^{N-1}\sum_{x_2=0}^{N-1}\psi(x_1,x_2,t)\,|x_1\rangle\,|x_2\rangle$, using $2n$ qubits with n qubits per spatial dimension.

Fourier Basis and Diagonal Kinetic Evolution. Let F denote the (unitary) discrete Fourier transform in each dimension. In the momentum/Fourier basis, $|\tilde{\psi}(t)\rangle = (F \otimes F)\,|\psi(t)\rangle$, the kinetic

operator becomes diagonal: $U_{\mathrm{kin}}(t) = \exp\Big(-i\,t\,\frac{1}{2}\hat{k}^2\Big)$, $\hat{k}^2\,|k_1\rangle\,|k_2\rangle = \lambda(k_1,k_2)\,|k_1\rangle\,|k_2\rangle$, where $\lambda(k_1,k_2)$ is the discrete Laplacian spectrum (proportional to $k_1^2+k_2^2$ under standard conventions). Thus, a baseline digital simulator applies $|\psi(t)\rangle \approx (F^\dagger\otimes F^\dagger)\,U_{\mathrm{kin}}(t)\,(F\otimes F)\,|\psi(0)\rangle$, and extends to $V\neq 0$ via operator splitting/Trotterization Berry et al. (2015a); Low & Chuang (2017); Childs et al. (2021); Berry et al. (2015b).

2.3 Task Operators and Low-Frequency Targets

Operator-First Objective. Rather than reconstructing the full field $\rho(x,t)$, ShadowFluid targets compact summaries such as low-frequency density structure or low-frequency energy. Let $\tilde{\rho}(q,t)$ denote a Fourier coefficient of the density at wavevector q (low-frequency when $\|q\|$ is small). A canonical identity connects $\tilde{\rho}(q,t)$ to second-order coherences of the Fourier amplitudes $\tilde{\rho}(q,t) \propto \sum_k b_{k+q}(t)\,b_k^*(t)$, where $b_k(t)$ denote the Fourier-mode coefficients of $\psi(\cdot,t)$. Low-frequency density statistics can therefore be represented by a restricted set of coherences rather than the full state.

Cutoff and Reduced Degrees of Freedom. Let $\mathcal{K}(K_0)$ be a low-frequency index set (e.g.,$\{k : \|k\|_\infty \le K_0\}$). Restricting the task operator set to indices within $\mathcal{K}(K_0)$ makes the reduced dimension M depend on K_0 rather than on N^2.

2.4 Shadow Hamiltonian Simulation

Heisenberg Evolution and Operator Closure. Following the shadow Hamiltonian framework Somma et al. (2025), let H be the Hamiltonian generating unitary evolution of a quantum state $\rho(t)$. For any operator O, the Heisenberg equation gives $\frac{d}{dt}\langle O\rangle_t = i\,\langle[H,O]\rangle_t, \qquad \langle O\rangle_t = \mathrm{Tr}\big(\rho(t)\,O\big)$. Consider an operator set $S=\{O_m\}_{m=1}^M$ spanning a linear subspace $\mathcal{S}=\mathrm{span}(S)$. A sufficient condition for closed dynamics of the expectations $\{\langle O_m\rangle_t\}$ is an invariance/closure property: $[H,O_m] = \sum_{m'=1}^M h_{mm'}\,O_{m'}, \qquad m=1,\dots,M$. Under (**??**), define the expectation vector $s(t)\in\mathbb{C}^M$ by $s_m(t)=\langle O_m\rangle_t$. Then it implies a linear ODE $\frac{d}{dt}s(t) = i\,H_S\,s(t), \qquad (H_S)_{mm'} = h_{mm'}$. When H_S is Hermitian (or can be made so with an appropriate operator basis choice), can be viewed as Schrödinger evolution in the reduced dimension M.

Shadow State. A shadow state is a normalized state $|\phi(t)\rangle\in\mathbb{C}^M$ whose amplitudes encode the expectation vector (up to a known linear map and normalization), $|\phi(t)\rangle \propto \sum_{m=1}^M s_m(t)\,|m\rangle$. Evolving $|\phi(t)\rangle$ under H_S advances all task operators in S simultaneously, replacing full Hilbert-space time-marching by evolution in dimension $M=|S|$.

2.5 Approximate Closure and Truncation

Exact closure may fail for realistic flow models (e.g.,with nontrivial V or nonlinear couplings). A standard strategy is to enforce approximate closure by projecting commutators back onto $\mathcal{S}$: $[H,O_m] = \sum_{m'=1}^M h_{mm'}\,O_{m'} + R_m$, where R_m is a residual outside $\mathcal{S}$. Defining a cutoff-controlled family of subspaces $\mathcal{S}(K_0)$ (e.g.,increasing low-frequency sets or bounded operator products) yields a tunable trade-off: larger K_0 increases M and cost, while reducing the magnitude of residuals and improving accuracy. ShadowFluid applies this idea by selecting S to target low-frequency density statistics and quantifying the accuracy-cost trade-off. We summarize key symbols: $\psi(x,t)$: wavefunction encoding flow state; $\rho(x,t)=|\psi(x,t)|^2$: density; F: discrete Fourier transform (implemented by QFT); $b_k(t)$: Fourier-mode coefficients; $S=\{O_m\}_{m=1}^M$: task operator set; $\mathcal{S}=\mathrm{span}(S)$: task subspace; H_S: induced shadow Hamiltonian/generator on $\mathbb{C}^M$; $\mathcal{K}(K_0)$: low-frequency index set by cutoff K_0.

3 Proposed ShadowFluid Method

ShadowFluid is an operator-first simulation framework: instead of propagating the full quantum state, it advances a task-selected set of operators under a low-dimensional shadow Hamiltonian and reconstructs the desired flow statistics from the evolved operators alone (Figure 1).

3.1 DICTIONARY DESIGN

Let $|\psi(t)\rangle$ be the quantum state encoding a discretized 2D flow (§ 2.2), evolving under a Hamiltonian H. Rather than reconstructing the full field $\rho(x,t) = |\psi(x,t)|^2$, we target a compact set of task outputs $y(t) = \big(f_1(t), \ldots, f_P(t)\big)$, where each $f_p(t)$ is an expectation of a task operator A_p: $f_p(t) = \langle A_p \rangle_t = \langle\psi(t)| A_p |\psi(t)\rangle$. Typical examples include low-frequency density Fourier coefficients $\tilde{\rho}(q,t)$ or low-frequency energy summaries. ShadowFluid constructs an operator set $S = \{O_m\}_{m=1}^{M}$ such that the task outputs are approximately linear functions of the expectation vector $s(t) = \big(\langle O_1\rangle_t, \ldots, \langle O_M\rangle_t\big)^\top$, and advances $s(t)$ using a reduced-order shadow generator.

Coherence-Based Operator Selection. In a QFT basis, write $|\psi(t)\rangle = \sum_k b_k(t)\,|k\rangle$. Low-frequency density coefficients satisfy $\tilde{\rho}(q,t) \propto \sum_k b_{k+q}(t)\, b_k^*(t)$, which is a sum of mode coherences, motivating the choice of S as a subset of rank-one operators $|k\rangle\langle k'|$ restricted by frequency cutoffs. We define two index sets: a task set $\mathcal{K}$ of low-frequency modes (controlled by a cutoff K_0) and a reference set $\mathcal{R}$. The multi-reference dictionary is $S(\mathcal{K},\mathcal{R}) = \big\{O_{k,r} = |k\rangle\langle r| : k \in \mathcal{K},\ r \in \mathcal{R}\big\}$, with expectations $Z_{k,r}(t) = \langle O_{k,r}\rangle_t = b_k(t)\, b_r^*(t)$, forming a $|\mathcal{K}| \times |\mathcal{R}|$ matrix $Z(t)$ (Figure 2). When $\mathcal{R} = \mathcal{K}$, $Z(t)$ is the low-frequency block of the density matrix. As a special case, fixing $\mathcal{R} = \{k_{\mathrm{ref}}\}$ recovers the single-reference dictionary used in our $V = 0$ prototype.

3.2 SHADOW DYNAMICS

When the Hamiltonian $H_0 = \sum_k E_k\,|k\rangle\langle k|$ is diagonal in the Fourier basis, the commutator closes exactly: $[H_0, O_{k,r}] = (E_k - E_r)\, O_{k,r}$. Any dictionary of coherences therefore satisfies the invariance property, and the shadow generator is diagonal. This regime serves as a sanity check: the shadow evolution matches the full low-pass evolution to machine precision.

When $H = H_0 + V$ with a nontrivial potential V, the Fourier-basis Hamiltonian acquires off-diagonal entries. For each coherence operator $O_{k,r} = |k\rangle\langle r|$, the commutator becomes $[H, O_{k,r}] = \sum_{k'} H_{k'k}\,|k'\rangle\langle r| - \sum_{r'} H_{rr'}\,|k\rangle\langle r'|$, which generates new operators along both the ket index ($k \to k'$) and the bra index ($r \to r'$). A single-reference dictionary (fixing $r = k_{\mathrm{ref}}$) generally does not close, because the second sum introduces operators $|k\rangle\langle r'|$ with $r' \neq k_{\mathrm{ref}}$.

Truncated Dynamics and Solution. The expectation $Z_{k,r}(t) = \mathrm{Tr}(\rho(t)\, O_{k,r})$ is the (r,k)-entry of the density matrix in Fourier basis; accordingly, $Z(t)$ constitutes a sub-block of the one-body reduced density matrix Coleman & Yukalov (2000). The von Neumann equation Breuer & Petruccione (2002) $\dot{\rho} = -i[H,\rho]$ implies $\dot{Z}_{k,r} = -i\sum_{k' \in \mathrm{all}} H_{kk'}\, Z_{k',r} + i\sum_{r' \in \mathrm{all}} H_{rr'}\, Z_{k,r'}$. Restricting both sums to the dictionary indices $k' \in \mathcal{K}$, $r' \in \mathcal{R}$ yields the truncated dynamics Quarteroni et al. (2015); Hesthaven et al. (2016) $\dot{Z} = -i\, H_{\mathcal{K}}\, Z + i\, Z\, H_{\mathcal{R}}$, where $H_{\mathcal{K}}$ and $H_{\mathcal{R}}$ are the sub-blocks of H restricted to $\mathcal{K}$ and $\mathcal{R}$, respectively. The solution is $Z(t) = U_{\mathcal{K}}(t)\, Z(0)\, U_{\mathcal{R}}(t)^\dagger$, where $U_{\mathcal{K}}(t) = e^{-iH_{\mathcal{K}}t}$ and $U_{\mathcal{R}}(t) = e^{-iH_{\mathcal{R}}t}$.

Reference-Set Construction. For structured potentials with sparse Fourier support, such as $V(x) = \sum_j \alpha_j \cos(q_j \cdot x)$, the off-diagonal coupling in Fourier space is local: H connects mode k to $k \pm q_j$. Starting from an initial reference r_0, we construct $\mathcal{R}$ by BFS on the coupling graph: $\mathcal{R}_0 = \{r_0\}$, $\mathcal{R}_{h+1} = \mathcal{R}_h \cup \big\{r \pm q_j : r \in \mathcal{R}_h,\ j = 1, \ldots, J\big\}$, iterated for $h = 0, \ldots, h_{\max} - 1$. For a single-component potential ($J = 1$), one hop gives $|\mathcal{R}| = 3$; two hops give $|\mathcal{R}| = 5$. The dictionary size is $M = |\mathcal{K}| \cdot |\mathcal{R}|$.

3.3 ERROR ANALYSIS

The truncation discards terms that couple $\mathcal{K}$ (resp. $\mathcal{R}$) to modes outside the dictionary. For each dictionary element $O_{k,r} \in S(\mathcal{K},\mathcal{R})$, the residual outside S has squared Frobenius norm $\|R_{k,r}\|_F^2 = \sum_{k' \notin \mathcal{K}} |H_{k'k}|^2 + \sum_{r' \notin \mathcal{R}} |H_{rr'}|^2$. This quantity is state-independent and can be computed before time evolution. We aggregate over the dictionary via an RMS measure: $\ell_{\mathrm{rms}} = \sqrt{\frac{1}{|\mathcal{K}||\mathcal{R}|} \sum_{k \in \mathcal{K},\, r \in \mathcal{R}} \|R_{k,r}\|_F^2}$.

Linear Error Bound. Let $L(t)$ denote the leakage matrix with entries $L_{k,r}(t) = \dot{Z}_{k,r}^{\mathrm{exact}}(t) - \dot{Z}_{k,r}^{\mathrm{trunc}}(t)$. Define $\Delta Z(t) = Z^{\mathrm{exact}}(t) - Z^{\mathrm{trunc}}(t)$; then ΔZ satisfies $\dot{\Delta Z} = -i\, H_{\mathcal{K}}\, \Delta Z +$

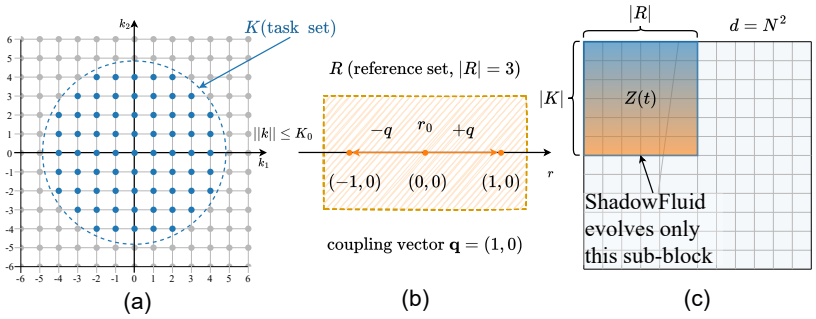


Figure 2: Dictionary construction via coupling graph. (a) Task set $\mathcal{K}(K_0)$. (b) Reference set $\mathcal{R}$ via 1-hop BFS. (c) Expectation matrix $Z(t)$ as a sub-block of the density matrix ρ.

Algorithm 1 ShadowFluid (Operator-First Shadow Evolution)

Require: Initial state $|\psi(0)\rangle$, Hamiltonian H, cutoff K_0
Ensure: Task outputs $y(t)$ and optional low-pass reconstruction

1: **Specify tasks:** Select target observables $y(t)$ (e.g., low-frequency density or energy summaries).
2: **Construct operator dictionary:** Define task index set $\mathcal{K}(K_0)$ and reference set $\mathcal{R}$.
3: **Initialize expectations:** Estimate the initial coherence matrix $Z(0) = \{b_k(0)\, b_r^*(0)\}_{k\in\mathcal{K},\, r\in\mathcal{R}}$ from the state $|\psi(0)\rangle$.
4: **Check truncation quality:** Compute the a priori leakage ℓ_{rms}; **if** ℓ_{rms} exceeds tolerance, **then** expand $\mathcal{K}$ or $\mathcal{R}$ and repeat Steps 2-4.
5: **Evolve reduced dynamics:** Propagate the truncated expectation matrix via $Z(t) = U_{\mathcal{K}}(t)\, Z(0)\, U_{\mathcal{R}}(t)^\dagger$, where $U_{\mathcal{K}}(t) = e^{-iH_{\mathcal{K}}t}$ and $U_{\mathcal{R}}(t) = e^{-iH_{\mathcal{R}}t}$.
6: **Reconstruct outputs:** Compute task observables $y(t)$ from $Z(t)$; reconstruct low-pass fields.

$i\,\Delta Z\, H_{\mathcal{R}} + L(t)$. Since the homogeneous part of the evolution consists of unitary conjugation (which preserves the Frobenius norm Higham (2008)), the variation-of-constants formula yields $\|\Delta Z(t)\|_F \le \int_0^t \|L(\tau)\|_F\, d\tau$.

Proof sketch. The variation-of-constants solution to with $\Delta Z(0) = 0$ is $\Delta Z(t) = \int_0^t U_{\mathcal{K}}(t-\tau)\, L(\tau)\, U_{\mathcal{R}}(t-\tau)^\dagger\, d\tau$. Taking the Frobenius norm and applying the triangle inequality, $\|\Delta Z(t)\|_F \le \int_0^t \|U_{\mathcal{K}}(t-\tau)\, L(\tau)\, U_{\mathcal{R}}(t-\tau)^\dagger\|_F\, d\tau$. Since $U_{\mathcal{K}}$ and $U_{\mathcal{R}}$ are unitary, $\|UAV^\dagger\|_F = \|A\|_F$, so each integrand reduces to $\|L(\tau)\|_F$. □

The bound grows linearly in t, not exponentially, because unitary conjugation preserves the Frobenius norm. Each entry $|L_{k,r}(t)| \le \|R_{k,r}\|_{\mathrm{op}}$ is bounded by operator norms independent of $\rho(t)$, so $\|L(t)\|_F$ can be upper-bounded by the a priori leakage, yielding a fully computable bound: $\|\Delta Z(t)\|_F \le t\, \ell_{\mathrm{rms}}\, \sqrt{|\mathcal{K}||\mathcal{R}|}$. For any linear task readout $y(t) = \mathrm{Tr}(W^\dagger Z(t))$: $|y_{\mathrm{exact}}(t) - y_{\mathrm{trunc}}(t)| \le \|W\|_F\, \|\Delta Z(t)\|_F$. We empirically validate that the a priori bound correlates with actual task error and that the linear-in-t behavior holds over the experimental time horizon (§ 4).

3.4 Pipeline and Reconstruction

For the block dictionary ($\mathcal{R} = \mathcal{K}$), the truncated evolution directly yields the low-frequency coefficients $b_k(t)$ for $k \in \mathcal{K}$. A low-pass field estimate follows as $\hat{\psi}_{\mathrm{LP}}(x,t) = \sum_{k\in\mathcal{K}(K_0)} \hat{b}_k(t)\, e^{ik\cdot x}$, $\hat{\rho}_{\mathrm{LP}}(x,t) = \left|\hat{\psi}_{\mathrm{LP}}(x,t)\right|^2$. For density-spectrum tasks, $\tilde{\rho}(q,t)$ can be assembled directly from $Z_{k+q,k}(t)$. Algorithm 1 summarizes the end-to-end pipeline.

The reduced dynamics operates on a $|\mathcal{K}| \times |\mathcal{R}|$ matrix Z whose size depends on the task bandwidth K_0 and the coupling structure of V, rather than on the full grid resolution N^2. A full-state simulation evolves $|\psi(t)\rangle$ in a Hilbert space of dimension $d = N^2$ and may require $O(N^2)$ measurements

for full-field characterization. ShadowFluid instead extracts only $O(|\mathcal{K}| \cdot |\mathcal{R}|)$ operator expectations Huang et al. (2020), thereby reducing the measurement complexity whenever the initial state is quantum-native (i.e.,produced by a quantum process and not efficiently representable classically).

4 EVALUATIONS

4.1 EXPERIMENTAL SETUP

Domain and Discretization. We use a periodic square domain $[0, 2\pi)^2$ on an $N \times N$ grid with $N = 2^{n_x}$, $n_x \in \{4, 5, 6\}$ ($N \in \{16, 32, 64\}$, Hilbert-space dimension $d = N^2 \in \{256, 1024, 4096\}$). Following Meng et al. (2024), the flow state is encoded as a two-component spinor $\psi = [\psi_1, \psi_2]^T$ so that the velocity field can carry nonzero vorticity, since a single-component wavefunction yields irrotational flow. Each component evolves independently under the same Hamiltonian H, and the density is $\rho = |\psi_1|^2 + |\psi_2|^2$. The framework of ShadowFluid applies per-component: we construct the dictionary and evolve $Z_j(t)$ for each ψ_j separately, then combine the results for task outputs. The initial condition is a Gaussian-vortex state generated by a deterministic seed-based procedure.

Potential. We consider single-component cosine potentials $V(x) = \alpha \cos(q \cdot x)$ with coupling vector $q = (1, 0)$ and coupling strength $\alpha \in \{0.1, 0.2, \ldots, 1.0\}$. This spans a $10\times$ range of coupling, from near-diagonal ($\alpha = 0.1$) to moderately off-diagonal ($\alpha = 1.0$).

Table 1: Configuration for V≠0 evaluation.

Parameter	Symbol	Values
Grid qubits/dim	n_x	$\{4, 5, 6\}$
Grid size	N	$\{16, 32, 64\}$
Hilbert dimension	$d = N^2$	$\{256, 1024, 4096\}$
Coupling strength	α	$\{0.1, 0.2, \ldots, 1.0\}$
Coupling vector	q	$(1, 0)$
Frequency cutoff	K_0	$1, 2, \ldots, 10$
Task-set size	$\|\mathcal{K}\|$	5-317
Reference hops	$h_{\max}$	1
Reference-set size	$\|\mathcal{R}\|$	3
Evolution time	t	$\{0.1, 0.2, \ldots, 1.0\}$

Dictionary design. The task set $\mathcal{K}(K_0)$ is a radial low-frequency set $\{k : \|k\| \le K_0\}$. We sweep $K_0 \in \{1, 2, \ldots, 10\}$. The reference set $\mathcal{R}$ is constructed by one-hop BFS closure from $r_0 = (0, 0)$ on the coupling graph, giving $|\mathcal{R}| = 3$ for the single-component potential. The dictionary size is $M = |\mathcal{K}| \cdot |\mathcal{R}|$.

Evolution. The full Hamiltonian $H = H_0 + V$ is constructed as a dense $d \times d$ matrix and eigendecomposed once per potential configuration. Full-state evolution computes $\psi(t) = e^{-iHt}\psi(0)$ via the eigenbasis. Truncated evolution computes $Z(t) = U_{\mathcal{K}}(t)\, Z(0)\, U_{\mathcal{R}}(t)^\dagger$ with $U_{\mathcal{K}} = e^{-iH_{\mathcal{K}}t}$ and $U_{\mathcal{R}} = e^{-iH_{\mathcal{R}}t}$.

Baselines. We compare ShadowFluid against two references: (i) *Full-state baseline*: exact evolution of the full d-dimensional state following the digital simulation pipeline of Meng et al. Meng et al. (2024) (QFT-based kinetic phase plus potential phase in the eigenbasis), yielding the ground-truth density ρ_{full} and Fourier coefficients $b_k^{\text{full}}(t)$; (ii) *Full low-pass*: the full-state coefficients truncated to $\mathcal{K}(K_0)$ at time t, yielding $\rho_{\text{full,LP}}$, the best possible low-frequency approximation under the same cutoff.

Table 2: V=0 sanity check: relative L^2 errors for three fluid observables ($N = 64$, $K_0 = 6$, $t = 0.3\pi$).

Observable	Shadow vs LP	Shadow vs Full	LP vs Full
Density ρ	1.85×10^{-16}	2.79×10^{-2}	2.79×10^{-2}
Vorticity ω	1.30×10^{-15}	2.12×10^{-1}	2.12×10^{-1}
Momentum $\|\mathbf{J}\|$	8.57×10^{-16}	1.91×10^{-1}	1.91×10^{-1}

Metrics. We report four error measures: $\varepsilon_\rho = \|\rho_{\text{trunc}} - \rho_{\text{full}}\|_2 / \|\rho_{\text{full}}\|_2$ (density error, truncated low-pass vs. full field); $\|\Delta Z\|_F = \|Z^{\text{exact}} - Z^{\text{trunc}}\|_F$ (Frobenius-norm error of the expectation matrix); $\varepsilon_{E_{\text{LP}}} = |E_{\text{LP}}^{\text{trunc}} - E_{\text{LP}}^{\text{full}}| / |E_{\text{LP}}^{\text{full}}|$ (relative error of the low-pass energy $E_{\text{LP}} = \sum_{k \in \mathcal{K}} (|b_1^k|^2 + |b_2^k|^2)$); and ℓ_{rms}, the state-independent a priori commutator leakage. Table 1 summarizes the parameter details.

4.2 EXPERIMENTAL RESULTS

RQ1: Sanity Check ($V = 0$). When $V = 0$, the Hamiltonian is diagonal in Fourier space, and the dictionary satisfies exact closure. In this regime, the shadow-evolved coefficients $\hat{b}_k(t)$ agree with the full-state low-pass coefficients to machine precision across all three fluid observables. Table 2

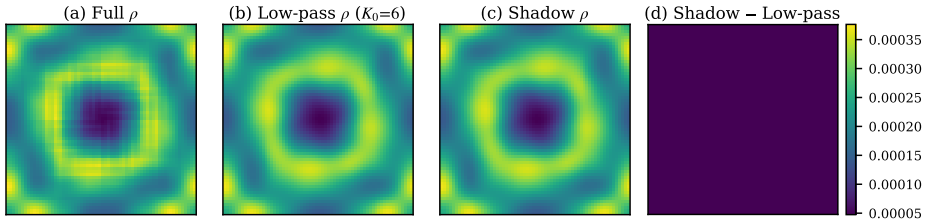


Figure 3: V=0 density field comparison ($N = 64$, $K_0 = 6$, $t = 0.3\pi$). (a) Full-state ρ; (b) low-pass ρ; (c) shadow ρ; (d) shadow minus low-pass residual. Panels (b) and (c) are visually identical; the residual in (d) is at machine precision ($\sim 10^{-16}$).

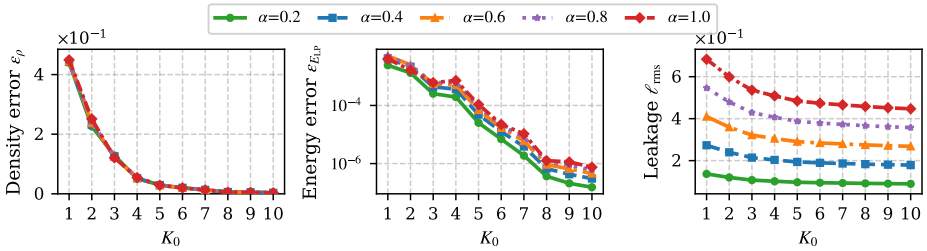


Figure 4: V≠0 errors versus cutoff K_0 at $t = 0.5$. Left: density error ε_ρ (truncated vs. full); center: task energy error $\varepsilon_{E_{\mathrm{LP}}}$; right: a priori leakage ℓ_{rms}. Different line styles and colors correspond to different coupling strengths $\alpha \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$.

reports the relative L^2 errors for density ρ, vorticity ω, and momentum magnitude $|\mathbf{J}|$ on the $N = 64$ grid with cutoff $K_0 = 6$ (113 retained modes out of 4096) at $t = 0.3\pi$. The shadow-vs-low-pass errors are at machine epsilon ($< 10^{-15}$), confirming exact agreement. The shadow-vs-full errors equal the low-pass-vs-full errors, showing that the only source of error in the $V \neq 0$ experiments is the off-diagonal truncation. Figure 3 visualizes the density field $\rho(x, t)$ at $t = 0.3\pi$ for the same configuration. Panels (a)-(c) show the full-state, low-pass, and shadow density fields on the 64×64 grid with $K_0 = 6$ (113 retained modes). The full-state field (a) contains fine-scale structure that is absent from the low-pass (b) and shadow (c) fields, which are visually indistinguishable from each other. Panel (d) confirms quantitatively that the shadow-minus-low-pass residual is at machine epsilon ($\sim 10^{-16}$), with colorbar values fourteen orders of magnitude below the density itself. This validates that the ShadowFluid coherence-based evolution is algebraically equivalent to direct low-pass truncation when $V = 0$.

RQ2: V≠0 Fidelity. Figure 4 plots the density error ε_ρ, task energy error $\varepsilon_{E_{\mathrm{LP}}}$, and a priori leakage ℓ_{rms} as functions of K_0 at $t = 0.5$ for five coupling strengths $\alpha \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$. Density error ε_ρ decreases monotonically from ∼0.22-0.25 at $K_0 = 2$ to $\sim 3.3 \times 10^{-3}$ at $K_0 = 10$, spanning nearly two orders of magnitude. Notably, ε_ρ is largely independent of α: even at $\alpha = 1.0$ (where $\ell_{\mathrm{rms}} \approx 0.45$), the density error at $K_0 = 10$ is 3.29×10^{-3}, comparable to the $\alpha = 0.2$ value of 3.27×10^{-3}. This indicates that the density error is dominated by the low-pass truncation itself, not by the dictionary approximation. Task energy error $\varepsilon_{E_{\mathrm{LP}}}$ is orders of magnitude smaller than ε_ρ, reaching $O(10^{-7})$ at $K_0 = 10$ for $\alpha = 1.0$. Leakage ℓ_{rms} scales linearly with α (as expected) and decreases slowly with K_0, since including more modes reduces the relative leakage per mode. Under $V \neq 0$ coupling, ShadowFluid achieves density errors below 1% at moderate cutoffs ($K_0 \geq 8$) across all tested coupling strengths spanning a $10\times$ range. The density error is controlled primarily by K_0 and is insensitive to α, indicating that the truncation introduces negligible additional error relative to simple low-pass filtering.

RQ3: Error Hierarchy. Figure 5 plots the error hierarchy at $\alpha = 0.5$, $t = 0.5$: the a priori bound $t \cdot \ell_{\mathrm{rms}} \sqrt{|\mathcal{K}||\mathcal{R}|}$, the actual $\|\Delta Z\|_F$, the density error ε_ρ, and the task energy error $\varepsilon_{E_{\mathrm{LP}}}$. The a priori bound always exceeds the actual $\|\Delta Z\|_F$. At $K_0 = 5$, for instance, the bound is 1.21×10^{-1} while the actual $\|\Delta Z\|_F$ is 7.73×10^{-3}, a gap of roughly one order of magnitude. A clear hierarchy holds at every K_0: bound $\geq \|\Delta Z\|_F \geq \varepsilon_\rho \geq \varepsilon_{E_{\mathrm{LP}}}$. The downstream task error is consistently 10^2-$10^4\times$ smaller than the operator-level error, confirming that low-frequency tasks are robust to the truncation-induced perturbation in Z. Because the bound is state-independent and computable before evolution, it serves as a practical quality indicator: if the bound is unacceptable, one can expand $\mathcal{K}$ or $\mathcal{R}$ before running the simulation. The a priori bound is valid and conservative. The

hierarchy bound $\geq \|\Delta Z\|_F \geq \varepsilon_\rho \geq \varepsilon_{E_{\mathrm{LP}}}$ holds uniformly, and task error is orders of magnitude below operator-level error.

RQ4: Temporal Stability. Figure 6 tracks the density error and $\|\Delta Z\|_F$ over time at $\alpha = 0.5$ for several representative cutoffs. $\|\Delta Z\|_F$ grows approximately linearly from $t = 0.1$ to $t = 1.0$: at $K_0 = 10$, it increases from 1.58×10^{-3} to 1.48×10^{-2}, roughly a $9\times$ increase over a $10\times$ increase in t. This confirms the linear-in-t prediction and rules out exponential error growth. In contrast, density error ε_ρ remains approximately flat over time (e.g.,$K_0 = 8$: from 6.96×10^{-3} at $t = 0.1$ to 5.73×10^{-3} at $t = 1.0$), because ε_ρ is dominated by the fixed low-pass truncation and the additional truncation error grows slowly. Task energy error remains consistently small ($< 10^{-6}$ for $K_0 \geq 8$) across all time points. $\|\Delta Z\|_F$ grows linearly in t as predicted by the unitary-preserving bound, with no exponential blowup. Density error is dominated by low-pass truncation and stays stable over time.

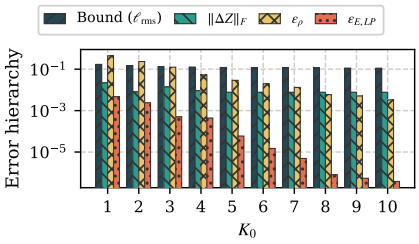


Figure 5: Error hierarchy at $\alpha = 0.5$, $t = 0.5$. The a priori bound is always valid: bound $\geq$ actual $\|\Delta Z\|_F \geq$ density error $\geq$ task error.

RQ5: Effect of Coupling Strength. Figure 7 shows how errors vary with coupling strength α at fixed $t = 0.5$ for representative cutoffs $K_0 \in \{3, 6, 10\}$. Leakage ℓ_{rms} scales linearly with α: at $K_0 = 10$, ℓ_{rms} goes from 0.045 ($\alpha = 0.1$) to 0.447 ($\alpha = 1.0$), a factor of $10\times$ for a $10\times$ increase in α. This is expected because $V(x) = \alpha \cos(q \cdot x)$ contributes off-diagonal entries proportional to $\alpha/2$. $\|\Delta Z\|_F$ also scales approximately linearly with α: at $K_0 = 10$, from 1.54×10^{-3} to 1.61×10^{-2}. Meanwhile, density error ε_ρ is nearly independent of α: at $K_0 = 10$, it ranges only from 3.27×10^{-3} to 3.29×10^{-3} across the full $10\times$ range, confirming that the density error reflects the low-pass truncation (dependent on K_0 alone) rather than the off-diagonal coupling. Leakage and operator-level error scale linearly with coupling strength, consistent with the perturbation scaling of V. Density error is insensitive to α, confirming that the dictionary approximation error remains a small perturbation on top of the dominant low-pass truncation.

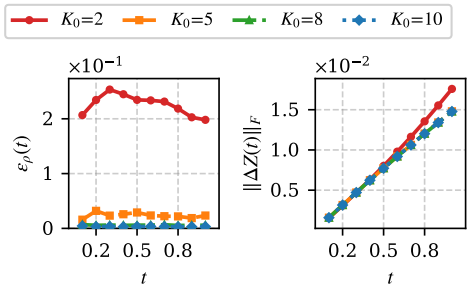


Figure 6: Temporal stability at $\alpha = 0.5$. Left: density error $\varepsilon_\rho(t)$; right: Frobenius error $\|\Delta Z(t)\|_F$. $\|\Delta Z\|_F$ grows approximately linearly in t as predicted, while density error remains largely flat.

RQ6: Multi-Component Potential. RQ1-RQ5 use single-component potentials $V(x) = \alpha \cos(q \cdot x)$ with a single coupling vector. We now test ShadowFluid on multi-component potentials $V(x) = \sum_{j=1}^{J} \alpha_j \cos(q_j \cdot x)$, where $J \in \{2, 4, 6, 8, 10\}$ Fourier components are drawn randomly with $\alpha_j \sim 0.1 \times \mathrm{Uniform}(0.5, 1.5)$ and q_j sampled uniformly from $\|q\|_\infty \leq 3$ (deterministic seed). The total coupling strength ranges from $\alpha_{\mathrm{tot}} \approx 0.26$ ($J = 2$) to ≈ 1.05 ($J = 10$). The denser coupling graph increases the reference set: one-hop BFS closure from $r_0 = (0, 0)$ yields $|\mathcal{R}| = 5$ ($J = 2$) up to $|\mathcal{R}| = 21$ ($J = 10$), compared with $|\mathcal{R}| = 3$ for the single-component case. Figure 8 plots the density error, task energy error, and $\|\Delta Z\|_F$ versus K_0 at $t = 0.5$ for $J \in \{2, 4, 6, 8, 10\}$, with the single-component Tier 1 baseline ($\alpha = 0.5$) shown as a dashed reference. Density error ε_ρ follows the same monotonic decrease with K_0 as in the single-component case: from $\sim$0.23-0.25 at $K_0 = 2$ to $\sim 3.2 \times 10^{-3}$-3.3×10^{-3} at $K_0 = 10$. The values are nearly identical across all J, confirming that low-pass truncation, not coupling complexity, controls ε_ρ.

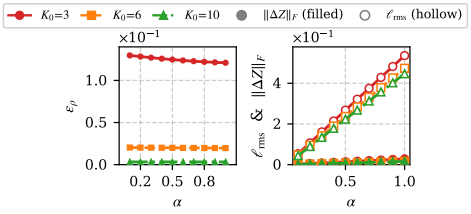


Figure 7: Effect of coupling strength at $t = 0.5$. Left: density error ε_ρ vs. α; right: leakage ℓ_{rms} (dashed) and $\|\Delta Z\|_F$ (solid) vs. α. Leakage and $\|\Delta Z\|_F$ scale linearly with α, while density error is nearly α-independent.

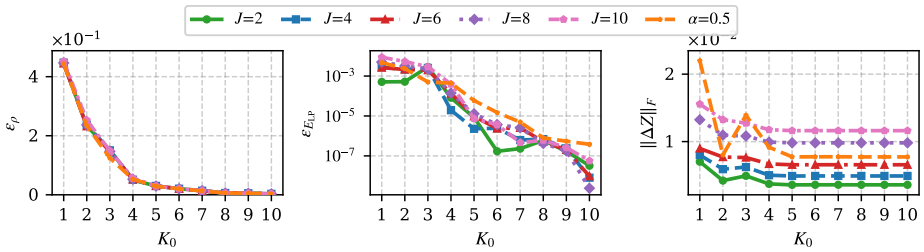


Figure 8: Multi-component potential errors versus cutoff K_0 at $t = 0.5$. Left: density error ε_ρ; center: task energy error $\varepsilon_{E_{\mathrm{LP}}}$; right: Frobenius error $\|\Delta Z\|_F$. Solid curves: $J \in \{2, 4, 6, 8, 10\}$ with increasing $|\mathcal{R}|$; dashed: single-component baseline ($\alpha = 0.5$).

$\|\Delta Z\|_F$ does increase with J: at $K_0 = 10$, $\|\Delta Z\|_F \approx 3.6 \times 10^{-3}$ ($J = 2$) versus 1.2×10^{-2} ($J = 10$), roughly proportional to α_{tot}. For $K_0 \geq 6$, however, $\|\Delta Z\|_F$ plateaus, indicating that the truncation error is saturated by the leakage from the fixed reference set rather than growing with $|\mathcal{K}|$. Task energy error $\varepsilon_{E_{\mathrm{LP}}}$ reaches $O(10^{-8})$-$O(10^{-9})$ at $K_0 = 10$ for all J values, orders of magnitude below the operator-level error. Leakage ℓ_{rms} ranges from 0.11 ($J = 2$) to 0.21 ($J = 10$), increasing with the number of coupling directions.

5 Related Work

Quantum Simulation of Fluid Dynamics. Quantum algorithms for PDEs and fluid-like dynamics have explored Schrödinger-flow formulations, QFT-based spectral methods, kinetic equations, and variational quantum simulation Harrow et al. (2009); Costa et al. (2019); Liu et al. (2021); Babbush et al. (2023); Bharadwaj & Sreenivasan (2023); Steijl & Barakos (2018); Oz et al. (2022); Budinski (2021); Succi et al. (2015); Childs & Liu (2020); Engel et al. (2019); Childs et al. (2017); Yuan et al. (2019); Cerezo et al. (2021).

Shadow Hamiltonian Simulation. Shadow Hamiltonian simulation Somma et al. (2025) shows that operator expectations obey reduced linear dynamics when the operator span is invariant under commutation with the Hamiltonian. ShadowFluid adapts this principle to quantum fluid dynamics and extends it to approximate closure, where finite dictionaries incur leakage.

Reduced-Order and Operator-Theoretic Models. Classical reduced-order models and Koopman-based methods compress high-dimensional dynamics using projection or data-driven operator representations Quarteroni et al. (2015); Benner et al. (2015); Lubich (2008); Mezić (2005); Williams et al. (2015); Brunton & Kutz (2022). Neural latent-dynamics methods provide learned alternatives Chen et al. (2018); Yin et al. (2021). In contrast, ShadowFluid constructs its reduced space directly from Hamiltonian coupling structure and derives linear-in-time error control from unitary reduced dynamics. Classical shadows reduce the measurement cost of estimating many observables Huang et al. (2020); Aaronson (2018); Huang et al. (2022), while quantum-inspired algorithms exploit classical low-rank structure Tang (2019).

Physics-Informed and Neural PDE Solvers. Physics-informed neural networks, neural operators, graph-based simulators, and ML-accelerated CFD learn surrogate dynamics from data or physical constraints Raissi et al. (2019); Karniadakis et al. (2021); Jia et al. (2021); Li et al. (2020); Bonev et al. (2023); Brandstetter et al. (2022); Sanchez-Gonzalez et al. (2020); Tompson et al. (2017); Kochkov et al. (2021); Wang et al. (2020).

6 Conclusion

In this paper, we present ShadowFluid, a task-aligned operator-first framework that replaces full-state quantum evolution with reduced dynamics over a coherence dictionary. The method provides a state-independent leakage criterion and a linear-in-time Frobenius error bound, enabling a priori control of the accuracy-dimension trade-off. Experiments on 2D Schrödinger-flow benchmarks show stable reduced dynamics and sub-1% density error at moderate cutoffs.

AI USE STATEMENT

Generative AI tools were used for language polishing and for adapting the manuscript's LaTeX formatting to the ICLR 2027 submission template. The authors are responsible for the final content of the paper.

REFERENCES

Scott Aaronson. Shadow tomography of quantum states. In *Proceedings of the 50th annual ACM SIGACT symposium on theory of computing*, pp. 325–338, 2018.

Ryan Babbush, Dominic W Berry, Robin Kothari, Rolando D Somma, and Nathan Wiebe. Exponential quantum speedup in simulating coupled classical oscillators. *Physical Review X*, 13(4): 041041, 2023.

Peter Benner, Serkan Gugercin, and Karen Willcox. A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM review*, 57(4):483–531, 2015.

Dominic W Berry, Andrew M Childs, Richard Cleve, Robin Kothari, and Rolando D Somma. Simulating hamiltonian dynamics with a truncated taylor series. *Physical review letters*, 114(9): 090502, 2015a.

Dominic W Berry, Andrew M Childs, and Robin Kothari. Hamiltonian simulation with nearly optimal dependence on all parameters. In *2015 IEEE 56th annual symposium on foundations of computer science*, pp. 792–809. IEEE, 2015b.

Sachin S Bharadwaj and Katepalli R Sreenivasan. Hybrid quantum algorithms for flow problems. *Proceedings of the National Academy of Sciences*, 120(49):e2311014120, 2023.

Boris Bonev, Thorsten Kurth, Christian Hundt, Jaideep Pathak, Maximilian Baust, Karthik Kashinath, and Anima Anandkumar. Spherical fourier neural operators: Learning stable dynamics on the sphere. In *International conference on machine learning*, pp. 2806–2823. PMLR, 2023.

Johannes Brandstetter, Daniel Worrall, and Max Welling. Message passing neural pde solvers. *arXiv preprint arXiv:2202.03376*, 2022.

Heinz-Peter Breuer and Francesco Petruccione. *The theory of open quantum systems*. OUP Oxford, 2002.

Steven L Brunton and J Nathan Kutz. *Data-driven science and engineering: Machine learning, dynamical systems, and control*. Cambridge University Press, 2022.

Ljubomir Budinski. Quantum algorithm for the advection–diffusion equation simulated with the lattice boltzmann method. *Quantum Information Processing*, 20(2), 2021.

Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, et al. Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644, 2021.

Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.

Andrew M Childs and Jin-Peng Liu. Quantum spectral methods for differential equations. *Communications in Mathematical Physics*, 375(2):1427–1457, 2020.

Andrew M Childs, Robin Kothari, and Rolando D Somma. Quantum algorithm for systems of linear equations with exponentially improved dependence on precision. *SIAM Journal on Computing*, 46(6):1920–1950, 2017.

Andrew M Childs, Yuan Su, Minh C Tran, Nathan Wiebe, and Shuchen Zhu. Theory of trotter error with commutator scaling. *Physical Review X*, 11(1):011020, 2021.

Albert John Coleman and Vyacheslav I Yukalov. *Reduced density matrices: Coulson's challenge*, volume 72. Springer Science & Business Media, 2000.

Pedro CS Costa, Stephen Jordan, and Aaron Ostrander. Quantum algorithm for simulating the wave equation. *Physical Review A*, 99(1):012323, 2019.

Alexander Engel, Graeme Smith, and Scott E Parker. Quantum algorithm for the vlasov equation. *Physical Review A*, 100(6):062315, 2019.

Iulia M Georgescu, Sahel Ashhab, and Franco Nori. Quantum simulation. *Reviews of Modern Physics*, 86(1):153–185, 2014.

Aram W Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical review letters*, 103(15):150502, 2009.

Jan S Hesthaven, Gianluigi Rozza, Benjamin Stamm, et al. *Certified reduced basis methods for parametrized partial differential equations*, volume 590. Springer, 2016.

Nicholas J Higham. *Functions of matrices: theory and computation*. SIAM, 2008.

Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020.

Hsin-Yuan Huang, Richard Kueng, Giacomo Torlai, Victor V Albert, and John Preskill. Provably efficient machine learning for quantum many-body problems. *Science*, 377(6613):eabk3333, 2022.

Xiaowei Jia, Jared Willard, Anuj Karpatne, Jordan S Read, Jacob A Zwart, Michael Steinbach, and Vipin Kumar. Physics-guided machine learning for scientific discovery: An application in simulating lake temperature profiles. *ACM/IMS Transactions on Data Science*, 2(3):1–26, 2021.

George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

Dmitrii Kochkov, Jamie A Smith, Ayya Alieva, Qing Wang, Michael P Brenner, and Stephan Hoyer. Machine learning–accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences*, 118(21):e2101784118, 2021.

Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.

Jin-Peng Liu, Herman Øie Kolden, Hari K Krovi, Nuno F Loureiro, Konstantina Trivisa, and Andrew M Childs. Efficient quantum algorithm for dissipative nonlinear differential equations. *Proceedings of the National Academy of Sciences*, 118(35):e2026805118, 2021.

Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996.

Guang Hao Low and Isaac L Chuang. Optimal hamiltonian simulation by quantum signal processing. *Physical review letters*, 118(1):010501, 2017.

Christian Lubich. *From quantum to classical molecular dynamics: reduced models and numerical analysis*. 2008.

Erwin Madelung. Quantum theory in hydrodynamical form. *z. Phys*, 40:322, 1927.

Zhaoyuan Meng, Jiarun Zhong, Shibo Xu, Ke Wang, Jiachen Chen, Feitong Jin, Xuhao Zhu, Yu Gao, Yaozu Wu, Chuanyu Zhang, et al. Simulating unsteady flows on a superconducting quantum processor. *Communications Physics*, 7(1):349, 2024.

Igor Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1):309–325, 2005.

Furkan Oz, Rohit KSS Vuppala, Kursat Kara, and Frank Gaitan. Solving burgers' equation with quantum computing. *Quantum Inf. Process.*, 21(1):30, 2022.

John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.

Alfio Quarteroni, Andrea Manzoni, and Federico Negri. *Reduced basis methods for partial differential equations: an introduction*, volume 92. Springer, 2015.

Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pp. 8459–8468. PMLR, 2020.

Rolando D Somma, Robbie King, Robin Kothari, Thomas E O'Brien, and Ryan Babbush. Shadow hamiltonian simulation. *Nature Communications*, 16(1):2690, 2025.

René Steijl and George N Barakos. Parallel evaluation of quantum algorithms for computational fluid dynamics. *Computers & Fluids*, 173:22–28, 2018.

Sauro Succi, François Fillion-Gourdeau, and Silvia Palpacelli. Quantum lattice boltzmann is a quantum walk. *EPJ Quantum Technology*, 2(1):1–17, 2015.

Ewin Tang. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing*, pp. 217–228, 2019.

Jonathan Tompson, Kristofer Schlachter, Pablo Sprechmann, and Ken Perlin. Accelerating eulerian fluid simulation with convolutional networks. In *International conference on machine learning*, pp. 3424–3433. PMLR, 2017.

Rui Wang, Karthik Kashinath, Mustafa Mustafa, Adrian Albert, and Rose Yu. Towards physics-informed deep learning for turbulent flow prediction. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 1457–1466, 2020.

Matthew O Williams, Ioannis G Kevrekidis, and Clarence W Rowley. A data–driven approximation of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, 2015.

Yuan Yin, Ibrahim Ayed, Emmanuel de Bézenac, Nicolas Baskiotis, and Patrick Gallinari. Leads: Learning dynamical systems that generalize across environments. *Advances in Neural Information Processing Systems*, 34:7561–7573, 2021.

Xiao Yuan, Suguru Endo, Qi Zhao, Ying Li, and Simon C Benjamin. Theory of variational quantum simulation. *Quantum*, 3:191, 2019.