# Q-Residual Physics: Hamiltonian-Structured Quantum Residual Learning for Embodied Dynamics

**Fanqi Kong**[1] **Ke Shi**[3] **Yuchen Wang**[1] **Zhipeng Liu**[1]
**Xiaofei Yue**[2] **Zhaoxuan Li**[4] **Tingting Li**[1*] **Ziming Zhao**[1*]
[1]Zhejiang University
[2]Beijing Institute of Technology
[3]University of Science and Technology of China
[4]State Key Laboratory of Cyberspace Security Defense, IIE, CAS, Beijing, China

## Abstract

Embodied agents in contact-rich environments must model dynamics affected by frictional uncertainty, contact discontinuities, compliance, actuator imperfections, perception noise, and mass or inertia mismatch. residual dynamics learning methods correct nominal physics predictions with unconstrained neural residuals, which often generalize poorly under unseen contact regimes, changing physical parameters, long-horizon rollouts, and few-shot adaptation. We propose **Q-Residual Physics**, a Hamiltonian-structured quantum residual learning framework for embodied dynamics. The method represents residual errors through a quantum residual layer whose qubits correspond to physically interpretable modes, including contact impulse, friction, slip, compliance, damping, perception error, and mass or inertia mismatch. A physical interaction graph induces a Hamiltonian residual structure, implemented with a Trotterized parameterized quantum circuit whose components capture mode activation, physical coupling, uncertain switching, dynamic mode exchange, and higher-order interactions. Quantum measurements are decoded to correct nominal physics predictions. Across embodied dynamics and manipulation benchmarks, Q-Residual Physics achieves stronger generalization than physics-only, neural residual, graph residual, and quantum-circuit baselines. Relative to the strongest neural residual baselines, it reduces 25-step rollout error by 25.5%, and reduces few-shot adaptation error by 25.2% with $10\%$ target-domain data.

## 1 Introduction

Embodied agents must predict and control complex physical dynamics in the presence of contact, friction, collisions, compliance, actuator imperfections, perception noise, and sim-to-real discrepancies. These factors introduce substantial deviations from idealized physical models: contact events lead to discontinuous state transitions, friction induces complex stick-slip behaviors, compliance and soft deformation produce nonlinear force responses, and imperfect sensing or inaccurate mass and inertia estimates further distort the observed dynamics. Explicit physics models provide a strong structural prior by encoding fundamental properties such as rigid-body motion, joint limits, contact constraints, and energy consistency. However, their idealized assumptions often fail to capture the full complexity of real-world interactions, particularly the subtle and coupled effects arising from frictional transitions, imperfect contact geometry, deformation, and parameter uncertainty. Residual dynamics learning provides a natural approach to addressing these limitations by combining a nominal next-state predictor [1] with a learned residual correction [2], thereby retaining the structural advantages of physics-based models while compensating for unmodeled dynamics.

*Co-corresponding authors: Tingting Li (litt2020@zju.edu.cn) and Ziming Zhao (zhaoziming@zju.edu.cn).

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

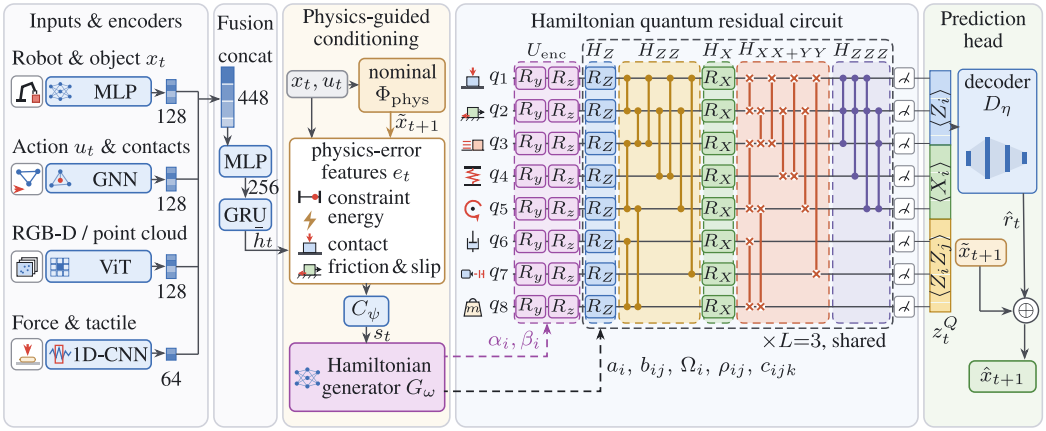


Figure 1: Q-Residual Physics overview

Despite its effectiveness, existing residual learning approaches typically employ generic neural architectures, including MLPs, latent dynamics models [3], graph neural networks [4], recurrent models [5], and transformers [6], to directly approximate the discrepancy between nominal predictions and real observations. While these architectures offer considerable expressive capacity, they generally treat the residual as an unconstrained black-box function, without explicitly modeling the underlying physical factors responsible for prediction errors. This can lead to entangled representations of distinct physical effects, making the learned corrections difficult to interpret and potentially limiting their generalization to unseen contact configurations, varying friction coefficients, different object masses, long-horizon rollouts, and few-shot adaptation scenarios. Our key observation is that residual errors in contact-rich environments are rarely arbitrary; rather, they often emerge from the interactions among a relatively structured set of physical residual modes, including contact impulses, friction, slip, compliance, perception uncertainty, and mass or inertia mismatch. These modes exhibit distinct physical characteristics while also influencing one another throughout the interaction process. Motivated by this observation, we argue that residual dynamics should be modeled as structured latent dynamics over interacting physical modes, rather than as a generic neural function, enabling more physically grounded representations of model discrepancies and potentially improving generalization and adaptation in complex embodied environments.

We propose *Q-Residual Physics*, a Hamiltonian-structured quantum residual learning framework. The method preserves the nominal physics model and replaces the black-box neural residual with a structured quantum layer. Physics-error features from the state, action, nominal prediction, contact indicators, constraint violations, energy and frictional features, and optional history or sensory encodings condition a Hamiltonian generator. Each qubit represents an interpretable residual mode, and a physical residual-mode graph specifies mode interactions. Figure 1 summarizes the architecture. Multimodal encoders and a GRU produce the history feature $h_t$, which joins the state, action, and nominal prediction $\tilde{x}_{t+1}$ in the physics-error features $e_t$. The conditioning network $C_\psi$ and Hamiltonian generator $G_\omega$ produce encoding angles and Hamiltonian coefficients for an 8-qubit circuit with three Trotter blocks. Its measurement vector $z_t^Q$ is decoded into $\hat{r}_t$ and integrated with $\tilde{x}_{t+1}$ to obtain $\hat{x}_{t+1}$.

The induced Hamiltonian is implemented by a Trotterized parameterized quantum circuit whose terms encode mode activation, pairwise physical coupling, uncertain switching, dynamic mode exchange, and sparse higher-order interactions. Quantum measurements form a structured residual representation, which a lightweight decoder maps to a nominal prediction correction. Learning is therefore constrained to a meaningful mode-interaction space rather than an arbitrary neural vector. This framework is a physics-informed inductive bias, not a claim of quantum hardware speedup. Unlike generic parameterized quantum circuits [7] or quantum-kernel hybrids [8], our topology and Hamiltonian terms are derived from residual-mode structure, supporting long-horizon stability, out-of-distribution generalization, and data-efficient adaptation.

We evaluate Q-Residual Physics on *OmniPush* [9], *ManiSkill2* [10], *robomimic* [11], *D4RL* [12], and *Physion* [13]. Across real and simulated benchmarks, it improves rollout accuracy, contact transition prediction, out-of-distribution generalization, and few-shot adaptation. Ablations over Hamiltonian

terms, topology, qubit number, Trotter depth, decoder capacity, and residual-mode dictionaries show the gains come from Hamiltonian-structured residual construction rather than a hand-crafted circuit.

In summary, this paper makes the following contributions:

- We introduce a Hamiltonian-structured quantum residual representation over interpretable physical modes and map their interaction graphs to Trotterized quantum circuits.
- We develop a physics-guided learning framework combining nominal prediction, physics-error conditioning, lightweight decoding, rollout training, consistency losses, and uncertainty objectives.
- We validate improved prediction, generalization, and transfer on five embodied dynamics benchmarks against neural, graph, generic quantum, and random circuit baselines, with ablations isolating the role of the Hamiltonian structure.

## 2 Preliminaries

### 2.1 Residual Dynamics Learning

We consider an embodied system with state $x_t \in \mathcal{X}$ and action $u_t \in \mathcal{U}$. The state may include proprioception, object pose, velocity, contact indicators, force or tactile readings, and optional visual or depth observations. Dynamics learning predicts $x_{t+1}$ or future rollout states. In contact-rich settings, contact modes, frictional transitions, compliance, perception uncertainty, and actuator imperfections make the transition function discontinuous and partially observed.

Model-based robot learning [1, 14, 15] uses dynamics predictors for control; residual learning [2] combines a nominal predictor with a learned correction. Let $\Phi_{\text{phys}}$ denote an explicit or approximate model encoding prior knowledge such as rigid-body motion, joint limits, contact response, friction approximations, and energy or constraint consistency:

$$\tilde{x}_{t+1} = \Phi_{\text{phys}}(x_t, u_t). \tag{1}$$

Since $\Phi_{\text{phys}}$ is imperfect, the true transition is modeled as

$$x_{t+1} = \text{Integrate}\big(\tilde{x}_{t+1}, r^*(x_t, u_t)\big), \tag{2}$$

where $r^*$ is the discrepancy between nominal and true dynamics. The integration reduces to addition in Euclidean spaces and applies the residual in the proper representation for rotations, joint limits, or contact constraints. A parametric residual model $r_\theta$ gives

$$\hat{x}_{t+1} = \text{Integrate}\big(\Phi_{\text{phys}}(x_t, u_t), \hat{r}_t\big), \qquad \hat{r}_t = r_\theta(x_t, u_t). \tag{3}$$

Learned predictors draw on architectures such as residual networks [16], MLPs, graph neural networks [17], recurrent models [18], and transformers [6] to parameterize $r_\theta$. Such residuals are expressive but unconstrained, and can learn spurious correlations, become unstable in long rollouts, or fail to transfer under changes in friction, mass, geometry, contact stiffness, or sensing. This motivates residual representations tied to physical error sources.

Residual models are trained for one-step prediction and rollout consistency. Given horizon $K$, the corrected model predicts

$$\hat{x}_{t+k} = F_\theta(\hat{x}_{t+k-1}, u_{t+k-1}), \qquad k = 1, \ldots, K, \tag{4}$$

where $F_\theta$ is the full corrected dynamics model. Since local residual errors accumulate, the residual should improve both local accuracy and long-horizon stability.

### 2.2 Structured Residual Modes

We assume residual errors decompose into physically meaningful modes, such as contact impulse error, frictional uncertainty, slip, compliance, rotation, damping, perception error, and mass or inertia mismatch. These modes interact, for example friction can trigger slip, slip can induce rotation, contact impulse can depend on compliance, and perception error can shift contact geometry, echoing interaction networks [19] and graph-based simulators [4, 20].

We define the residual mode set and interaction graph as

$$\mathcal{M} = \{m_1, m_2, \ldots, m_N\}, \tag{5}$$

$$\mathcal{G} = (\mathcal{M}, E, T), \tag{6}$$

where each $m_i$ is an interpretable residual source, $E$ contains pairwise couplings, and $T$ contains sparse higher-order interactions. An edge $(i, j) \in E$ denotes coupled modes such as friction and slip, while a tuple $(i, j, k) \in T$ denotes a joint effect such as contact, friction, and rotation. The graph is an inductive bias, not a full simulator.

In Q-Residual Physics, each mode maps to one qubit. Qubits do not encode robot state dimensions; they represent latent physical error factors. In the 8-qubit contact-rich manipulation instance, the modes are contact impulse, friction, slip, compliance, rotation, joint damping, depth perception error, and mass or inertia mismatch. Other tasks may use different dictionaries while preserving the same mapping from residual modes to a Hamiltonian-structured circuit.

The residual-mode graph induces a Hamiltonian over qubits. Let $X_i$, $Y_i$, and $Z_i$ denote Pauli operators on qubit $i$, and let $|0\rangle_i$ and $|1\rangle_i$ denote the inactive and active states of mode $m_i$. We use

$$\begin{aligned} H_{\text{res}}(e_t) = &\sum_i a_i(e_t) Z_i + \sum_{(i,j)\in E} b_{ij}(e_t) Z_i Z_j + \sum_i \Omega_i(e_t) X_i \\ &+ \sum_{(i,j)\in E} \rho_{ij}(e_t)(X_i X_j + Y_i Y_j) + \sum_{(i,j,k)\in T} c_{ijk}(e_t) Z_i Z_j Z_k. \end{aligned} \tag{7}$$

Here $e_t$ contains physics-error features from the state, action, nominal prediction, constraints, energy, contact, friction, and optional history or sensory encodings. A conditioning network generates $a_i$, $b_{ij}$, $\Omega_i$, $\rho_{ij}$, and $c_{ijk}$ from $e_t$, making the Hamiltonian state- and task-dependent while preserving residual-mode structure.

The $Z_i$ terms activate individual modes, $Z_i Z_j$ models pairwise coupling, $X_i$ models uncertain switching, $X_i X_j + Y_i Y_j$ models dynamic exchange, which transfers activation between coupled modes while conserving their total $Z_i + Z_j$, as in friction-to-slip or slip-to-rotation transfer, and $Z_i Z_j Z_k$ captures sparse higher-order interactions. We implement this Hamiltonian with a Trotterized circuit that encodes $e_t$ into single-qubit rotations, applies repeated Hamiltonian blocks, and measures expectations such as $\langle Z_i \rangle$, $\langle X_i \rangle$, and $\langle Z_i Z_j \rangle$. These measurements form a structured residual representation decoded into $\hat{r}_t$.

# 3 Q-Residual Physics

## 3.1 Overview and Physics-Error Conditioning

Q-Residual Physics keeps explicit physics as the nominal predictor and uses a Hamiltonian-structured quantum residual layer to model the remaining error. Given state $x_t$, action $u_t$, and optional sensory observations $o_t$, a standard encoder summarizes the observable context, while the nominal physics module produces the first-order prediction. We collect the quantities used for residual correction as

$$\begin{aligned} &\bar{h}_t = E_\phi(x_t, u_t, o_t), \qquad \tilde{x}_{t+1} = \Phi_{\text{phys}}(x_t, u_t), \\ &e_t = \Big[x_t,\, u_t,\, \tilde{x}_{t+1},\, g_{\text{constraint}},\, g_{\text{energy}},\, g_{\text{contact}},\, g_{\text{friction}},\, \bar{h}_t\Big], \\ &s_t = C_\psi(e_t), \qquad \hat{x}_{t+1} = \text{Integrate}\Big(\Phi_{\text{phys}}(x_t, u_t), D_\eta(z_t^Q)\Big). \end{aligned} \tag{8}$$

Here $g_{\text{constraint}}$, $g_{\text{energy}}$, $g_{\text{contact}}$, and $g_{\text{friction}}$ summarize constraint margins, energy consistency, contact information, and friction or slip cues. The encoder $E_\phi$ (Figure 1) combines an MLP for robot and object state, a graph network for actions and contact relations, a frozen ViT for RGB-D or point-cloud observations, and a 1D-CNN for force and tactile signals; modalities absent from a benchmark are omitted. Their 128-, 128-, 128-, and 64-dimensional outputs are concatenated into a 448-dimensional feature, projected to 256 dimensions by an MLP, and summarized over the observation history by a GRU to give $\bar{h}_t$. The key design choice is that the quantum residual layer is conditioned on physics-error features rather than raw observations alone, so it corrects nominal physics instead of relearning the entire dynamics. In additive Euclidean state spaces, the final line reduces to $\hat{x}_{t+1} = \Phi_{\text{phys}}(x_t, u_t) + D_\eta(z_t^Q)$; for rotations, articulated systems, or constrained states, the residual is applied through the integration operator.

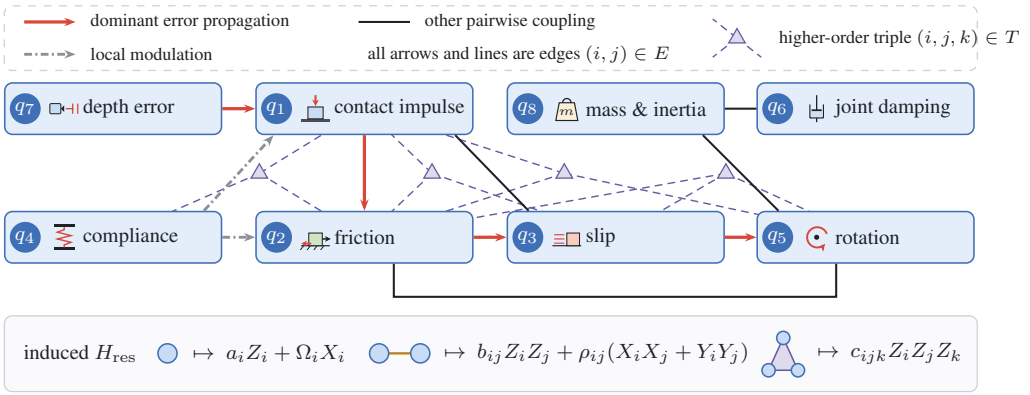


Figure 2: Residual-mode interaction graph

### 3.2 Hamiltonian Quantum Residual Layer

The core of Q-Residual Physics is a residual layer whose qubits correspond to interpretable residual modes rather than raw state dimensions. In the default $N = 8$ instantiation, qubits $q_1, \ldots, q_8$ represent contact impulse, friction, slip, compliance, rotation, joint damping, depth perception error, and mass/inertia mismatch, respectively. The graph $\mathcal{G} = (\mathcal{M}, E, T)$ of Figure 2 restricts interactions to physically related modes. Its edge set

$$E = \{(1,2), (1,3), (1,4), (1,7), (2,3), (2,4), (2,5), (3,5), (5,8), (6,8)\} \tag{9}$$

contains the dominant depth–contact–friction–slip–rotation chain, the modulation of contact and friction by compliance, the direct contact–slip and friction–rotation couplings, and the effect of mass and inertia mismatch on damping and rotation. The triple set

$$T = \{(1,2,3), (1,2,4), (1,2,5), (2,3,5)\} \tag{10}$$

captures joint effects of contact or friction with a downstream mode: the friction-cone condition for slip onset, compliance-dependent friction, friction-induced torque, and slip-induced rotation. The graph can be specified from domain knowledge or partially learned with sparsity regularization, and the mode dictionary can vary across tasks. Figure 2 distinguishes dominant error propagation (red arrows), local compliance modulation (grey arrows), and other edges in $E$ (black lines). Arrows indicate physical propagation; the induced couplings are symmetric. Triangles denote triples in $T$, and the bottom panel maps nodes, edges, and triples to Hamiltonian terms. Given the residual context $s_t$, a Hamiltonian generator emits the data-encoding angles and the coefficients of $H_{\text{res}}$ in Equation (7), $\{\alpha_i, \beta_i, a_i, b_{ij}, \Omega_i, \rho_{ij}, c_{ijk}\} = G_\omega(s_t)$, $4N + 2|E| + |T| = 56$ numbers in the default instance. To control the evolution scale, each generated coefficient is smoothly bounded: $v = v_{\max} \tanh(\hat{v})$ for $v \in \{a_i, b_{ij}, \Omega_i, \rho_{ij}, c_{ijk}\}$ with $v_{\max} = 1$, and the encoding angles are bounded in the same way by $\pi$. This parameterization remains differentiable end to end.

### 3.3 Trotterized Circuit Implementation and Training

The Hamiltonian residual dynamics are implemented using a Trotterized parameterized quantum circuit related to quantum feature maps [21], data re-uploading [22], and trainable circuit models [7]. The encoding unitary is $U_{\text{enc}} = \prod_i R_z^{(i)}(\beta_i) R_y^{(i)}(\alpha_i)$, which acts on $|0\rangle^{\otimes N}$ with $R_y$ first, so that $\langle Z_i \rangle = \cos \alpha_i$ sets the prior activation of mode $i$ and $\beta_i$ sets the relative phases on which the exchange terms act. The block in Figure 3 applies, from left to right, mode activation on every qubit, pairwise coupling on the $|E| = 10$ edges, uncertain switching on every qubit, dynamic exchange on the same edges, and higher-order interaction on the $|T| = 4$ triples. We repeat this block $L = 3$ times after $U_{\text{enc}}$ by default. Writing $\Delta t = 1/L$ for the evolution step, the complete circuit is

$$U_{\text{res}}(e_t) = \left[ \prod_{\ell=1}^{L} U^{(\ell)} \right] U_{\text{enc}}, \qquad U^{(\ell)} = U_{ZZZ} U_{XX+YY} U_X U_{ZZ} U_Z. \tag{11}$$

Using $R_P(\vartheta) = \exp(-\mathrm{i}\vartheta P/2)$ for a Pauli string $P$, the factors $U_Z, U_{ZZ}, U_X$, and $U_{ZZZ}$ are products of $R_{Z_i}(2\Delta t\, a_i)$, $R_{Z_i Z_j}(2\Delta t\, b_{ij})$, $R_{X_i}(2\Delta t\, \Omega_i)$, and $R_{Z_i Z_j Z_k}(2\Delta t\, c_{ijk})$, respectively, over their

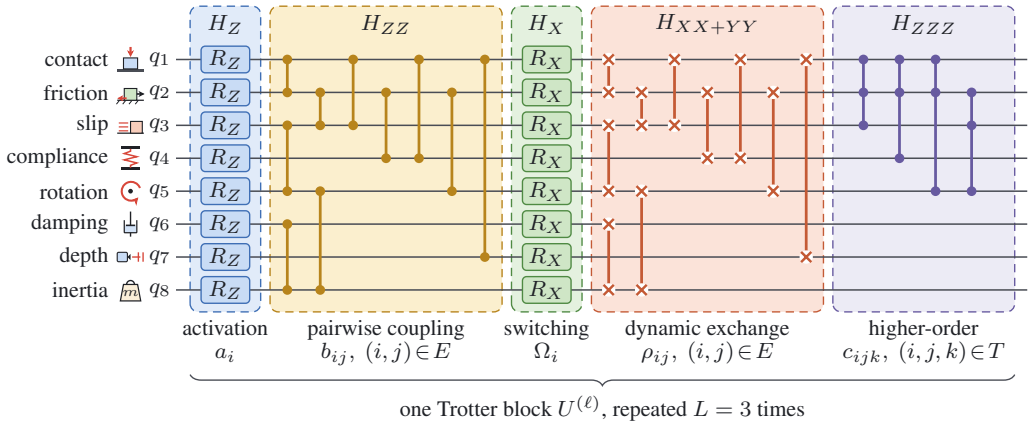


Figure 3: Trotterized residual block

corresponding modes, edges, or triples. The exchange factor uses $R_{X_iX_j}(2\Delta t\,\rho_{ij})R_{Y_iY_j}(2\Delta t\,\rho_{ij})$ on each edge, which implements $\exp\big(-\mathrm{i}\Delta t\,\rho_{ij}(X_iX_j + Y_iY_j)\big)$ exactly because $X_iX_j$ and $Y_iY_j$ commute. The coefficients are shared across blocks, so the circuit is a first-order Trotterization of $\exp(-\mathrm{i}H_{\mathrm{res}})$ and $L$ trades circuit depth against Trotter accuracy. After quantum evolution, we measure expectation values that summarize individual mode activation, uncertain switching, and pairwise coupling:

$$z_t^Q = \Big[\{\langle Z_i\rangle\}_{i=1}^N, \{\langle X_i\rangle\}_{i=1}^N, \{\langle Z_iZ_j\rangle\}_{(i,j)\in E}\Big] \in \mathbb{R}^{2N+|E|}, \tag{12}$$

which is 26-dimensional in the default instance. Sparse triple expectations $\{\langle Z_iZ_jZ_k\rangle\}_{(i,j,k)\in T}$ can be appended when higher-order residual effects are important. The measurement vector is decoded into a residual correction by a lightweight decoder $D_\eta$. For uncertainty estimation, the decoder instead outputs a residual mean and log variance. We deliberately keep $D_\eta$ small so that the main residual representation is carried by the Hamiltonian-structured quantum bottleneck rather than by a high-capacity neural shortcut (Table 5). Training uses multi-step rollout supervision with physics-consistency regularization. Let $F_\theta$ denote the corrected dynamics model, recursively applied as $\hat{x}_{t+k} = F_\theta(\hat{x}_{t+k-1}, u_{t+k-1})$. The objective is

$$\mathcal{L} = \sum_{k=1}^{K} \|\hat{x}_{t+k} - x_{t+k}\|_2^2 + \lambda_c\mathcal{L}_{\mathrm{constraint}} + \lambda_E\mathcal{L}_{\mathrm{energy}} + \lambda_s\mathcal{L}_{\mathrm{stability}} + \lambda_{\mathrm{nll}}\mathcal{L}_{\mathrm{uncertainty}}. \tag{13}$$

The constraint loss penalizes joint-limit violations, object penetration, and friction-cone violations. The energy loss penalizes physically implausible energy growth beyond estimated input work, while the stability loss regularizes rollout drift and residual smoothness. The uncertainty term is used only when the decoder predicts a residual distribution; otherwise $\lambda_{\mathrm{nll}} = 0$. These terms encourage physically plausible corrections over long rollouts.

## 4 Experiments

We evaluate Q-Residual Physics from four perspectives. First, we test whether the proposed Hamiltonian-structured quantum residual layer improves standard dynamics prediction in contact-rich embodied settings. Second, we evaluate out-of-distribution physical generalization under changes in friction, mass, inertia, contact stiffness, and object geometry. Third, we study data-efficient adaptation in few-shot and sim-to-real settings. Fourth, we perform ablations over Hamiltonian terms, circuit topology, qubit number, Trotter depth, decoder capacity, and residual-mode dictionaries to verify that the improvement comes from the proposed residual-mode Hamiltonian construction rather than from a particular hand-crafted circuit.

### 4.1 Setup and Baselines

**Benchmarks.** We evaluate on five embodied dynamics benchmarks: *OmniPush* [9], *ManiSkill2* [10], *robomimic* [11], *D4RL* [12], and *Physion* [13], covering real pushing, manipulation, demonstration

dynamics, adaptation, and physical prediction. Models use standard training splits and are evaluated on held-out trajectories.

**Inputs and prediction targets.** State inputs include robot proprioception, object pose and velocity, actions, and available contact indicators. RGB-D or point-cloud observations are processed by a lightweight frozen visual encoder when used. Contact relations are constructed from proximity, estimated contact normals, relative velocities, and penetration or distance features. Prediction targets are planar object motion for pushing, robot–object states for manipulation, future states for adaptation, and object trajectories for physical prediction; translation and rotation are evaluated separately when possible, with joint and end-effector states used for articulated systems.

**Models and baselines.** All methods use the same benchmark-specific observations, nominal physics model, encoder, train/test split, and rollout protocol. Our default Q-Residual Physics model uses an 8-qubit residual-mode circuit with three Trotter blocks and a two-layer MLP decoder. We compare with *Physics Only*, *MLP Residual*, *GNN Residual* [17, 4], *GRU Residual* [18, 5], *Generic PQC* [7, 23], *Random Sparse PQC*, and *Ising Quantum Residual*.

**Metrics.** We report normalized one-step, 10-step, and 25-step prediction errors, contact transition accuracy, slip F1, energy and constraint violations, OOD generalization error, and few-shot adaptation error. Lower is better except for contact accuracy, slip F1, and manipulation success rate.

### 4.2 Dynamics Prediction Results

**Contact-rich dynamics prediction.** Table 1 reports one-step and multi-step prediction results across the five benchmarks. Q-Residual Physics achieves the lowest average one-step error, 10-step rollout error, and 25-step rollout error. Averaged across benchmarks, Q-Residual Physics obtains errors of $0.038$, $0.098$, and $0.161$, compared with $0.054$, $0.136$, and $0.216$ for the strongest neural residual baseline at the same horizons. These correspond to relative reductions of $29.6\%$, $27.9\%$, and $25.5\%$, respectively. The improvement remains large at long horizons, supporting the hypothesis that a Hamiltonian-structured residual representation is especially useful when local prediction errors accumulate over time.

Table 1: Dynamics Prediction Error Results

| Metric | Method | *OmniPush* | *ManiSkill2* | *robomimic* | *D4RL* | *Physion* | Avg. |
|---|---|---|---|---|---|---|---|
| One-step error ↓ | Physics Only | 0.091 | 0.112 | 0.104 | 0.097 | 0.126 | 0.106 |
| | MLP Residual | 0.052 | 0.061 | 0.058 | 0.055 | 0.069 | 0.059 |
| | GNN Residual | 0.046 | 0.055 | 0.053 | 0.051 | 0.063 | 0.054 |
| | GRU Residual | 0.047 | 0.057 | 0.051 | 0.049 | 0.064 | 0.054 |
| | Generic PQC | 0.056 | 0.065 | 0.061 | 0.058 | 0.074 | 0.063 |
| | Random Sparse PQC | 0.054 | 0.063 | 0.060 | 0.057 | 0.071 | 0.061 |
| | Ising Quantum Residual | 0.043 | 0.052 | 0.050 | 0.048 | 0.060 | 0.051 |
| | **Q-Residual Physics** | **0.032**$^{-0.011}$ | **0.041**$^{-0.011}$ | **0.039**$^{-0.011}$ | **0.037**$^{-0.011}$ | **0.043**$^{-0.017}$ | **0.038**$^{-0.013}$ |
| 10-step rollout ↓ | Physics Only | 0.241 | 0.284 | 0.262 | 0.251 | 0.309 | 0.269 |
| | MLP Residual | 0.132 | 0.158 | 0.149 | 0.143 | 0.171 | 0.151 |
| | GNN Residual | 0.119 | 0.146 | 0.137 | 0.132 | 0.159 | 0.139 |
| | GRU Residual | 0.121 | 0.142 | 0.134 | 0.129 | 0.156 | 0.136 |
| | Generic PQC | 0.145 | 0.169 | 0.160 | 0.153 | 0.184 | 0.162 |
| | Random Sparse PQC | 0.139 | 0.164 | 0.155 | 0.150 | 0.179 | 0.157 |
| | Ising Quantum Residual | 0.112 | 0.134 | 0.128 | 0.124 | 0.148 | 0.129 |
| | **Q-Residual Physics** | **0.083**$^{-0.029}$ | **0.103**$^{-0.031}$ | **0.097**$^{-0.031}$ | **0.092**$^{-0.032}$ | **0.114**$^{-0.034}$ | **0.098**$^{-0.031}$ |
| 25-step rollout ↓ | Physics Only | 0.417 | 0.502 | 0.471 | 0.438 | 0.552 | 0.476 |
| | MLP Residual | 0.219 | 0.264 | 0.251 | 0.237 | 0.294 | 0.253 |
| | GNN Residual | 0.187 | 0.231 | 0.218 | 0.207 | 0.253 | 0.219 |
| | GRU Residual | 0.193 | 0.226 | 0.211 | 0.201 | 0.249 | 0.216 |
| | Generic PQC | 0.242 | 0.286 | 0.269 | 0.253 | 0.315 | 0.273 |
| | Random Sparse PQC | 0.231 | 0.277 | 0.260 | 0.247 | 0.303 | 0.264 |
| | Ising Quantum Residual | 0.176 | 0.214 | 0.202 | 0.193 | 0.237 | 0.204 |
| | **Q-Residual Physics** | **0.132**$^{-0.044}$ | **0.165**$^{-0.049}$ | **0.156**$^{-0.046}$ | **0.147**$^{-0.046}$ | **0.205**$^{-0.032}$ | **0.161**$^{-0.043}$ |

**Contact and slip prediction.** Table 2 reports contact transition accuracy, slip prediction F1, and physical consistency metrics. Q-Residual Physics improves contact accuracy from $74.6\%$ for GNN Residual to $83.9\%$, and improves slip F1 from $0.722$ to $0.801$. It also exceeds the Ising Quantum

Residual baseline, which reaches $76.1\%$ contact accuracy and $0.735$ slip F1. These gains are consistent with the design of the residual Hamiltonian: the $X$ terms model uncertain switching, while the $XX + YY$ exchange terms model transitions between friction, slip, and rotation modes.

Table 2: Contact and Slip Prediction

| Method | Contact Acc. ↑ | Slip F1 ↑ | Energy Viol. ↓ | Constraint Viol. ↓ |
|---|---|---|---|---|
| Physics Only | 61.8 | 0.512 | 0.184 | 0.147 |
| MLP Residual | 70.9 | 0.671 | 0.103 | 0.084 |
| GNN Residual | 74.6 | 0.722 | 0.089 | 0.071 |
| GRU Residual | 73.8 | 0.706 | 0.092 | 0.074 |
| Generic PQC | 69.4 | 0.642 | 0.112 | 0.090 |
| Random Sparse PQC | 70.1 | 0.653 | 0.108 | 0.087 |
| Ising Quantum Residual | 76.1 | 0.735 | 0.081 | 0.066 |
| **Q-Residual Physics** | **83.9** | **0.801** | **0.057** | **0.044** |

### 4.3 Generalization and Transfer

**Out-of-distribution physical generalization.** We evaluate OOD generalization by training on source physical regimes and testing on held-out regimes with shifted friction, mass, inertia, contact stiffness, and object geometry. Figure 4(a) reports errors for each physical shift and their average. Q-Residual Physics consistently outperforms all baselines, reaching average OOD error $0.124$ compared with $0.159$ for GNN Residual and $0.166$ for GRU Residual, a $22.0\%$ reduction relative to the strongest neural residual baseline.

**Sim-to-real transfer.** We train on simulation and adapt to real or more realistic target dynamics using limited real-domain trajectories. Figure 4(b) reports real-domain rollout error and manipulation success rate. Q-Residual Physics reduces error from $0.184$ for GNN Residual to $0.137$ and improves success rate from $61.8\%$ to $70.6\%$. In the same transfer evaluation, contact accuracy improves from $71.5\%$ to $79.6\%$, and slip F1 from $0.689$ to $0.773$. These results support transfer across mismatches in friction, compliance, and contact geometry.

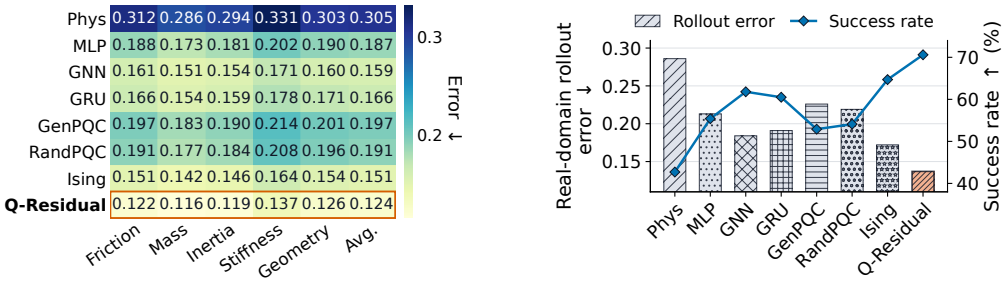


(a) OOD shifts

(b) Sim-to-real transfer

Figure 4: Generalization and transfer

### 4.4 Ablations and Efficiency

**Hamiltonian term ablation.** Table 4 studies the contribution of different Hamiltonian components. Removing the exchange term $XX+YY$ increases 10-step rollout error from $0.098$ to $0.119$, increases 25-step rollout error from $0.161$ to $0.192$, and reduces slip F1 from $0.801$ to $0.756$. Removing the higher-order $ZZZ$ terms increases 10-step rollout error to $0.113$ and 25-step rollout error to $0.184$, affecting collision-centric and non-center pushing scenarios. The Ising-only variant, which keeps only $Z$ and $ZZ$ terms, performs better than generic PQC but worse than the full model. This suggests that static mode activation and pairwise coupling are useful, but uncertain switching, dynamic exchange, and higher-order interactions are necessary for the best performance.

Table 3: Ablation on circuit size and depth.

| Ablation | Setting | Error ↓ | | | Cost ↓ |
|---|---|---|---|---|---|
| | | 10-step | 25-step | OOD | |
| Topology (Params.) | Physics Graph | **0.098** | **0.161** | **0.124** | 1.00× |
| | Learned Sparse | 0.101 | 0.166 | 0.127 | 1.08× |
| | Fully Connected | 0.112 | 0.184 | 0.143 | 1.42× |
| | Ring | 0.126 | 0.207 | 0.156 | 0.93× |
| | Chain | 0.132 | 0.216 | 0.164 | 0.90× |
| | Random Sparse | 0.157 | 0.264 | 0.191 | 1.00× |
| Qubits (Train time) | $N = 4$ | 0.131 | 0.213 | 0.166 | 0.72× |
| | $N = 6$ | 0.109 | 0.178 | 0.137 | 0.86× |
| | $N = 8$ | **0.098** | **0.161** | **0.124** | 1.00× |
| | $N = 12$ | 0.096 | 0.163 | 0.129 | 1.54× |
| Trotter depth (Train time) | $L = 1$ | 0.124 | 0.205 | 0.153 | 0.71× |
| | $L = 2$ | 0.104 | 0.171 | 0.131 | 0.86× |
| | $L = 3$ | **0.098** | **0.161** | **0.124** | 1.00× |
| | $L = 4$ | 0.096 | 0.160 | 0.128 | 1.19× |

Table 4: Hamiltonian term ablation.

| Variant | 10 | 25 | Acc. | F1 |
|---|---|---|---|---|
| Full | **0.098** | **0.161** | **83.9** | **0.801** |
| w/o $XX + YY$ | 0.119 | 0.192 | 79.2 | 0.756 |
| w/o $ZZZ$ | 0.113 | 0.184 | 80.6 | 0.771 |
| w/o $X$ | 0.116 | 0.188 | 78.4 | 0.748 |
| $Z + ZZ$ only | 0.129 | 0.204 | 76.1 | 0.735 |
| Single-qubit | 0.148 | 0.238 | 71.3 | 0.681 |

Table 5: Decoder and mode ablation.

| Variant | 1 | 10 | OOD | FS |
|---|---|---|---|---|
| Full | 0.038 | **0.098** | **0.124** | **0.113** |
| Large dec. | 0.036 | 0.101 | 0.139 | 0.126 |
| Shortcut | **0.035** | 0.103 | 0.146 | 0.132 |
| No Q bott. | 0.054 | 0.139 | 0.159 | 0.151 |
| Random modes | 0.049 | 0.128 | 0.153 | 0.142 |
| w/o percept. | 0.041 | 0.107 | 0.134 | 0.121 |
| w/o compl. | 0.043 | 0.112 | 0.141 | 0.127 |
| w/o inertia | 0.042 | 0.110 | 0.138 | 0.124 |

**Circuit topology, qubit number, and Trotter depth.** Table 3 reports circuit topology, qubit number, and Trotter-depth ablations. The physics-induced residual-mode graph performs best, with 10-step error $0.098$, 25-step error $0.161$, and OOD error $0.124$. A learned sparse graph performs similarly with error $0.101$, but requires additional regularization and tuning. Fully connected circuits perform worse than the physics graph, likely due to over-parameterization and weaker inductive bias. For qubit number and Trotter depth, four qubits lack enough residual modes and reach 10-step error $0.131$. Increasing to six qubits improves the error to $0.109$. Increasing to 12 qubits slightly lowers the 10-step error to $0.096$ but worsens OOD error to $0.129$ and increases training time to $1.54\times$, suggesting possible overfitting or unnecessary residual-mode complexity.

**Decoder capacity and residual-mode dictionary.** Table 5 studies whether the gains come from the quantum residual bottleneck or from the residual decoder. A large decoder improves in-distribution one-step error but worsens OOD error, suggesting that excessive decoder capacity can bypass the structured quantum representation and overfit. A shortcut decoder that directly concatenates physics-error features with quantum measurements also reduces in-distribution error but hurts OOD performance. The best OOD results are obtained with a lightweight decoder, supporting our design choice. Replacing physically interpretable modes with random latent modes raises 10-step error from $0.098$ to $0.128$; removing perception, compliance, or inertia modes also increases OOD error, supporting the physical mode dictionary.

**Few-shot data efficiency.** Models are first trained on source domains and then fine-tuned using $1\%$, $5\%$, $10\%$, or $25\%$ of target-domain trajectories. Figure 5(a) reports 10-step rollout error at each data budget. Q-Residual Physics reaches errors of $0.176$, $0.132$, $0.113$, and $0.097$, respectively, compared with $0.244$, $0.178$, $0.151$, and $0.126$ for GNN Residual. At $10\%$ target data, the reduction relative to GNN Residual is $25.2\%$, supporting the data efficiency of the structured residual representation.

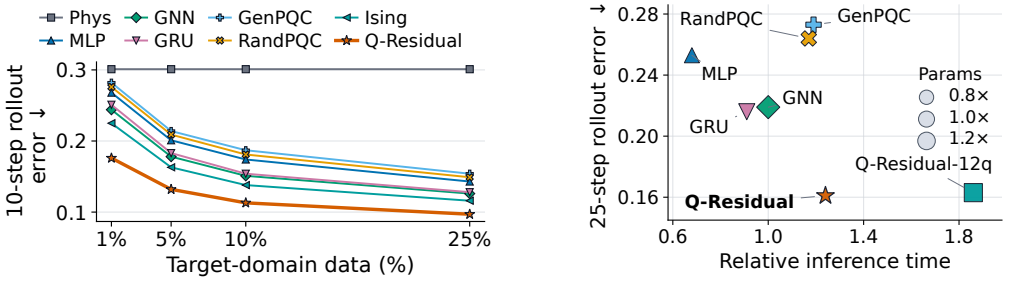


(a) Few-shot adaptation

(b) Accuracy–cost trade-off

Figure 5: Data and computational efficiency

**Computational cost.** Figure 5(b) compares average 25-step rollout error, inference time, and parameter count (marker size), with costs normalized to GNN Residual. The default 8-qubit model uses $0.93\times$ the parameters, $1.31\times$ training time, and $1.24\times$ inference time, while reducing error from $0.219$ to $0.161$. Quantum circuit simulation adds cost over MLP and GNN residual baselines but remains practical at this scale. The 12-qubit variant needs $1.22\times$ the parameters, $2.02\times$ training time, and $1.86\times$ inference time for error $0.163$, supporting the default 8-qubit design.

## 5 Related Work

**Residual and contact-rich dynamics.** Residual learning combines structured predictors with learned corrections in vision [16] and robot control [2]. Related model-based robot learning includes policy search [1, 14], model predictive control [24], latent dynamics [3, 25], neural dynamics for manipulation and control [26, 15], and recurrent world models [18, 5]. Contact-rich dynamics are commonly evaluated in pushing [9], manipulation [10, 11], offline adaptation [12], and physical prediction [13] benchmarks. Broader embodied environments, including RLBench [27], robosuite [28], SAPIEN [29], MuJoCo [30], Isaac Gym [31], SoftGym [32], and Meta-World [33], further emphasize physical structure. Few-shot and sim-to-real methods address distribution shifts through domain randomization [34], dynamics randomization [35], domain adaptation [36], robust policy optimization [37], meta-learning [38], meta-reinforcement learning [39], and simulation adaptation from real-world experience [40]. Complementary approaches include sequence modeling [6] and model-based offline policy optimization [41–43].

**Graph-based and physics-structured dynamics.** Graph and relational models encode interacting objects and contacts through interaction networks [19], visual interaction networks [44], neural relational inference [45], graph-network physics engines [17, 4], and mesh simulators [20]. Equivariant architectures [46–48] additionally encode geometric symmetries. Physics-informed models impose differential-equation structure [49, 50], Lagrangian dynamics [51, 52], Hamiltonian dynamics [53–55], and symplectic recurrence [56]. Unlike these approaches, Q-Residual Physics defines graph nodes as residual modes rather than scene objects or system states, and uses graph edges to induce Hamiltonian couplings among physical error factors.

**Quantum machine learning.** Quantum machine learning [57] includes feature maps [21], data re-uploading [22], parameterized circuits [7], and quantum-kernel hybrids [8]. Trainability depends on optimization geometry [58], barren plateaus [59], cost functions [60], and architecture [61]. PennyLane [23], TensorFlow Quantum [62], and TorchQuantum [63] support hybrid modeling; related studies address embedding search [64], transfer [65], quanvolution [66], capacity [67], and variational optimization [68]. QuFrame [69] combines encoding schemes, QuST [70] trains quantum neural networks against spatial and temporal noise, and QuFid [71] uses circuit graphs for adaptive noise-aware fidelity estimation. QCDeploy [72] uses graph contrastive learning and ranking for circuit deployment, while QuaMap [73] benchmarks backend-aware mapping and learning-based compiler evaluation. Our graph instead encodes physical residual interactions; generic and random circuit baselines isolate the role of Hamiltonian structure.

## 6 Conclusion and Limitations

We introduced *Q-Residual Physics*, a Hamiltonian-structured quantum residual learning framework for embodied dynamics. By modeling residual errors through interacting physical modes, it improves rollout accuracy, contact prediction, OOD generalization, sim-to-real transfer, and few-shot adaptation. Ablations validate the importance of Hamiltonian-structured residual modeling.

**Limitations.** The method relies on predefined residual modes and interaction graphs, potentially missing complex physical couplings. Hardware deployment and scalability remain constrained by noise, circuit depth, and trainability issues, while applicability is limited to tasks with physically structured residuals.

## Acknowledgments and Disclosure of Funding

This work was supported by the Ningbo Yongjiang Talent Programme (2025A-387-G), the Zhejiang Provincial Natural Science Foundation of China (Grant No. LQN26F020008), the CCF-Kuaishou Large Model Explorer Fund (No. CCF-KuaiShou 2025001), the CAAI-Ant Group Research Fund (2025CAAI-ANT-15), the Science and Technology Innovation Project of Ningbo Global Innovation Center, Zhejiang University (Grant No. S02J2604), and the China Postdoctoral Science Foundation under Grant No. GZB20260393.

## References

[1] Marc Peter Deisenroth and Carl Edward Rasmussen. PILCO: A model-based and data-efficient approach to policy search. In Lise Getoor and Tobias Scheffer, editors, *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pages 465–472. Omnipress, 2011. URL https://icml.cc/2011/papers/323_icmlpaper.pdf.

[2] Tobias Johannink, Shikhar Bahl, Ashvin Nair, Jianlan Luo, Avinash Kumar, Matthias Loskyll, Juan Aparicio Ojea, Eugen Solowjow, and Sergey Levine. Residual reinforcement learning for robot control. In *International Conference on Robotics and Automation, ICRA 2019, Montreal, QC, Canada, May 20-24, 2019*, pages 6023–6029. IEEE, 2019. doi: 10.1109/ICRA.2019.8794127. URL https://doi.org/10.1109/ICRA.2019.8794127.

[3] Manuel Watter, Jost Tobias Springenberg, Joschka Boedecker, and Martin A. Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2746–2754, 2015. URL https://proceedings.neurips.cc/paper/2015/hash/a1afc58c6ca9540d057299ec3016d726-Abstract.html.

[4] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter W. Battaglia. Learning to simulate complex physics with graph networks. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, Proceedings of Machine Learning Research, pages 8459–8468. PMLR, 2020. URL http://proceedings.mlr.press/v119/sanchez-gonzalez20a.html.

[5] Danijar Hafner, Timothy P. Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL https://openreview.net/forum?id=S1lOTC4tDS.

[6] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Michael Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 15084–15097, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/7f489f642a0ddb10272b5c31057f0663-Abstract.html.

[7] Marcello Benedetti, Erika Lloyd, and Stefan H. Sack. Parameterized quantum circuits as machine learning models. *CoRR*, abs/1906.07682, 2019. URL http://arxiv.org/abs/1906.07682.

[8] Daniel T. Chang. Parameterized quantum circuits with quantum kernels for machine learning: A hybrid quantum-classical approach. *CoRR*, abs/2209.14449, 2022. doi: 10.48550/ARXIV.2209.14449. URL https://doi.org/10.48550/arXiv.2209.14449.

[9] Maria Bauzá, Ferran Alet, Yen-Chen Lin, Tomás Lozano-Pérez, Leslie Pack Kaelbling, Phillip Isola, and Alberto Rodriguez. Omnipush: accurate, diverse, real-world dataset of pushing dynamics with RGB-D video. In *2019 IEEE/RSJ International Conference on Intelligent*

*Robots and Systems, IROS 2019, Macau, SAR, China, November 3-8, 2019*, pages 4265–4272. IEEE, 2019. doi: 10.1109/IROS40897.2019.8967920. URL https://doi.org/10.1109/IROS40897.2019.8967920.

[10] Jiayuan Gu, Fanbo Xiang, Xuanlin Li, Zhan Ling, Xiqiang Liu, Tongzhou Mu, Yihe Tang, Stone Tao, Xinyue Wei, Yunchao Yao, Xiaodi Yuan, Pengwei Xie, Zhiao Huang, Rui Chen, and Hao Su. Maniskill2: A unified benchmark for generalizable manipulation skills. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL https://openreview.net/forum?id=b_CQDy9vrD1.

[11] Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. In Aleksandra Faust, David Hsu, and Gerhard Neumann, editors, *Conference on Robot Learning, 8-11 November 2021, London, UK*, Proceedings of Machine Learning Research, pages 1678–1690. PMLR, 2021. URL https://proceedings.mlr.press/v164/mandlekar22a.html.

[12] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4RL: datasets for deep data-driven reinforcement learning. *CoRR*, abs/2004.07219, 2020. URL https://arxiv.org/abs/2004.07219.

[13] Daniel Bear, Elias Wang, Damian Mrowca, Felix J. Binder, Hsiao-Yu Tung, R. T. Pramod, Cameron Holdaway, Sirui Tao, Kevin A. Smith, Fan-Yun Sun, Fei-Fei Li, Nancy Kanwisher, Josh Tenenbaum, Dan Yamins, and Judith E. Fan. Physion: Evaluating physical prediction from vision in humans and machines. In Joaquin Vanschoren and Sai-Kit Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*, 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/d09bf41544a3365a46c9077ebb5e35c3-Abstract-round1.html.

[14] Sergey Levine and Vladlen Koltun. Guided policy search. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, JMLR Workshop and Conference Proceedings, pages 1–9. JMLR.org, 2013. URL http://proceedings.mlr.press/v28/levine13.html.

[15] Anusha Nagabandi, Gregory Kahn, Ronald S. Fearing, and Sergey Levine. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018*, pages 7559–7566. IEEE, 2018. doi: 10.1109/ICRA.2018.8463189. URL https://doi.org/10.1109/ICRA.2018.8463189.

[16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.90. URL https://doi.org/10.1109/CVPR.2016.90.

[17] Alvaro Sanchez-Gonzalez, Nicolas Heess, Jost Tobias Springenberg, Josh Merel, Martin A. Riedmiller, Raia Hadsell, and Peter W. Battaglia. Graph networks as learnable physics engines for inference and control. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, Proceedings of Machine Learning Research, pages 4467–4476. PMLR, 2018. URL http://proceedings.mlr.press/v80/sanchez-gonzalez18a.html.

[18] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 2455–2467, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/2de5d16682c3c35007e4e92982f1a2ba-Abstract.html.

[19] Peter W. Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, and Koray Kavukcuoglu. Interaction networks for learning about objects, relations and physics. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4502–4510, 2016. URL https://proceedings.neurips.cc/paper/2016/hash/3147da8ab4a0437c15ef51a5cc7f2dc4-Abstract.html.

[20] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W. Battaglia. Learning mesh-based simulation with graph networks. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL https://openreview.net/forum?id=roNqYL0_XP.

[21] Vojtech Havlícek, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nat.*, 567(7747):209–212, 2019. doi: 10.1038/S41586-019-0980-2. URL https://doi.org/10.1038/s41586-019-0980-2.

[22] Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I. Latorre. Data re-uploading for a universal quantum classifier. *Quantum*, 4:226, 2020. doi: 10.22331/Q-2020-02-06-226. URL https://doi.org/10.22331/q-2020-02-06-226.

[23] Ville Bergholm, Josh A. Izaac, Maria Schuld, Christian Gogolin, and Nathan Killoran. Pennylane: Automatic differentiation of hybrid quantum-classical computations. *CoRR*, abs/1811.04968, 2018. URL http://arxiv.org/abs/1811.04968.

[24] Ian Lenz, Ross A. Knepper, and Ashutosh Saxena. Deepmpc: Learning deep latent features for model predictive control. In Lydia E. Kavraki, David Hsu, and Jonas Buchli, editors, *Robotics: Science and Systems XI, Sapienza University of Rome, Rome, Italy, July 13-17, 2015*, 2015. doi: 10.15607/RSS.2015.XI.012. URL http://www.roboticsproceedings.org/rss11/p12.html.

[25] Danijar Hafner, Timothy P. Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, Proceedings of Machine Learning Research, pages 2555–2565. PMLR, 2019. URL http://proceedings.mlr.press/v97/hafner19a.html.

[26] Anusha Nagabandi, Kurt Konolige, Sergey Levine, and Vikash Kumar. Deep dynamics models for learning dexterous manipulation. In Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura, editors, *3rd Annual Conference on Robot Learning, CoRL 2019, Osaka, Japan, October 30 - November 1, 2019, Proceedings*, Proceedings of Machine Learning Research, pages 1101–1112. PMLR, 2019. URL http://proceedings.mlr.press/v100/nagabandi20a.html.

[27] Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J. Davison. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics Autom. Lett.*, 5(2):3019–3026, 2020. doi: 10.1109/LRA.2020.2974707. URL https://doi.org/10.1109/LRA.2020.2974707.

[28] Yuke Zhu, Josiah Wong, Ajay Mandlekar, and Roberto Martín-Martín. robosuite: A modular simulation framework and benchmark for robot learning. *CoRR*, abs/2009.12293, 2020. URL https://arxiv.org/abs/2009.12293.

[29] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, and Hao Su. SAPIEN: A simulated part-based interactive environment. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 11094–11104. Computer Vision Foundation / IEEE, 2020. doi: 10.1109/CVPR42600.2020.01111. URL https://openaccess.thecvf.com/content_CVPR_2020/html/Xiang_SAPIEN_A_SimulAted_Part-Based_Interactive_ENvironment_CVPR_2020_paper.html.

[30] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2012, Vilamoura, Algarve, Portugal, October 7-12, 2012*, pages 5026–5033. IEEE, 2012. doi: 10.1109/IROS.2012.6386109. URL https://doi.org/10.1109/IROS.2012.6386109.

[31] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac gym: High performance GPU based physics simulation for robot learning. In Joaquin Vanschoren and Sai-Kit Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*, 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/28dd2c7955ce926456240b2ff0100bde-Abstract-round2.html.

[32] Xingyu Lin, Yufei Wang, Jake Olkin, and David Held. Softgym: Benchmarking deep reinforcement learning for deformable object manipulation. In Jens Kober, Fabio Ramos, and Claire J. Tomlin, editors, *4th Conference on Robot Learning, CoRL 2020, 16-18 November 2020, Virtual Event / Cambridge, MA, USA*, Proceedings of Machine Learning Research, pages 432–448. PMLR, 2020. URL https://proceedings.mlr.press/v155/lin21a.html.

[33] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura, editors, *3rd Annual Conference on Robot Learning, CoRL 2019, Osaka, Japan, October 30 - November 1, 2019, Proceedings*, Proceedings of Machine Learning Research, pages 1094–1100. PMLR, 2019. URL http://proceedings.mlr.press/v100/yu20a.html.

[34] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2017, Vancouver, BC, Canada, September 24-28, 2017*, pages 23–30. IEEE, 2017. doi: 10.1109/IROS.2017.8202133. URL https://doi.org/10.1109/IROS.2017.8202133.

[35] Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018*, pages 1–8. IEEE, 2018. doi: 10.1109/ICRA.2018.8460528. URL https://doi.org/10.1109/ICRA.2018.8460528.

[36] Konstantinos Bousmalis, Alex Irpan, Paul Wohlhart, Yunfei Bai, Matthew Kelcey, Mrinal Kalakrishnan, Laura Downs, Julian Ibarz, Peter Pastor, Kurt Konolige, Sergey Levine, and Vincent Vanhoucke. Using simulation and domain adaptation to improve efficiency of deep robotic grasping. In *2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018*, pages 4243–4250. IEEE, 2018. doi: 10.1109/ICRA.2018.8460875. URL https://doi.org/10.1109/ICRA.2018.8460875.

[37] Aravind Rajeswaran, Sarvjeet Ghotra, Balaraman Ravindran, and Sergey Levine. Epopt: Learning robust neural network policies using model ensembles. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL https://openreview.net/forum?id=SyWvgP5el.

[38] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, Proceedings of Machine Learning Research, pages 1126–1135. PMLR, 2017. URL http://proceedings.mlr.press/v70/finn17a.html.

[39] Anusha Nagabandi, Ignasi Clavera, Simin Liu, Ronald S. Fearing, Pieter Abbeel, Sergey Levine, and Chelsea Finn. Learning to adapt in dynamic, real-world environments through meta-reinforcement learning. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL https://openreview.net/forum?id=HyztsoC5Y7.

[40] Yevgen Chebotar, Ankur Handa, Viktor Makoviychuk, Miles Macklin, Jan Issac, Nathan D. Ratliff, and Dieter Fox. Closing the sim-to-real loop: Adapting simulation randomization with real world experience. In *International Conference on Robotics and Automation, ICRA 2019, Montreal, QC, Canada, May 20-24, 2019*, pages 8973–8979. IEEE, 2019. doi: 10.1109/ICRA.2019.8793789. URL https://doi.org/10.1109/ICRA.2019.8793789.

[41] Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Y. Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma. MOPO: model-based offline policy optimization. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/a322852ce0df73e204b7e67cbbef0d0a-Abstract.html.

[42] Tianhe Yu, Aviral Kumar, Rafael Rafailov, Aravind Rajeswaran, Sergey Levine, and Chelsea Finn. COMBO: conservative offline model-based policy optimization. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 28954–28967, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/f29a179746902e331572c483c45e5086-Abstract.html.

[43] Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. Morel: Model-based offline reinforcement learning. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/f7efa4f864ae9b88d43527f4b14f750f-Abstract.html.

[44] Nicholas Watters, Daniel Zoran, Theophane Weber, Peter W. Battaglia, Razvan Pascanu, and Andrea Tacchetti. Visual interaction networks: Learning a physics simulator from video. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 4539–4547, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/8cbd005a556ccd4211ce43f309bc0eac-Abstract.html.

[45] Thomas N. Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard S. Zemel. Neural relational inference for interacting systems. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, Proceedings of Machine Learning Research, pages 2693–2702. PMLR, 2018. URL http://proceedings.mlr.press/v80/kipf18a.html.

[46] Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E(n) equivariant graph neural networks. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, Proceedings of Machine Learning Research, pages 9323–9332. PMLR, 2021. URL http://proceedings.mlr.press/v139/satorras21a.html.

[47] Fabian Fuchs, Daniel E. Worrall, Volker Fischer, and Max Welling. Se(3)-transformers: 3d roto-translation equivariant attention networks. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/15231a7ce4ba789d13b722cc5c955834-Abstract.html.

[48] Nathaniel Thomas, Tess E. Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation- and translation-equivariant neural networks for 3d point clouds. *CoRR*, abs/1802.08219, 2018. URL http://arxiv.org/abs/1802.08219.

[49] Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 6572–6583, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/69386f6bb1dfed68692a24c8686939b9-Abstract.html.

[50] 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. *J. Comput. Phys.*, 378:686–707, 2019. doi: 10.1016/J.JCP.2018.10.045. URL https://doi.org/10.1016/j.jcp.2018.10.045.

[51] Michael Lutter, Christian Ritter, and Jan Peters. Deep lagrangian networks: Using physics as model prior for deep learning. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL https://openreview.net/forum?id=BklHpjCqKm.

[52] Miles D. Cranmer, Sam Greydanus, Stephan Hoyer, Peter W. Battaglia, David N. Spergel, and Shirley Ho. Lagrangian neural networks. *CoRR*, abs/2003.04630, 2020. URL https://arxiv.org/abs/2003.04630.

[53] Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 15353–15363, 2019. URL https://proceedings.neurips.cc/paper/2019/hash/26cd8ecadce0d4efd6cc8a8725cbd1f8-Abstract.html.

[54] Yaofeng Desmond Zhong, Biswadip Dey, and Amit Chakraborty. Symplectic ode-net: Learning hamiltonian dynamics with control. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL https://openreview.net/forum?id=ryxmb1rKDS.

[55] Peter Toth, Danilo J. Rezende, Andrew Jaegle, Sébastien Racanière, Aleksandar Botev, and Irina Higgins. Hamiltonian generative networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL https://openreview.net/forum?id=HJenn6VFvB.

[56] Zhengdao Chen, Jianyu Zhang, Martín Arjovsky, and Léon Bottou. Symplectic recurrent neural networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL https://openreview.net/forum?id=BkgYPREtPr.

[57] Jacob D. Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nat.*, 549(7671):195–202, 2017. doi: 10.1038/NATURE23474. URL https://doi.org/10.1038/nature23474.

[58] James Stokes, Josh A. Izaac, Nathan Killoran, and Giuseppe Carleo. Quantum natural gradient. *Quantum*, 4:269, 2020. doi: 10.22331/Q-2020-05-25-269. URL https://doi.org/10.22331/q-2020-05-25-269.

[59] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *CoRR*, abs/1803.11173, 2018. URL http://arxiv.org/abs/1803.11173.

[60] Marco Cerezo, Akira Sone, Tyler J. Volkoff, Lukasz Cincio, and Patrick J. Coles. Cost-function-dependent barren plateaus in shallow quantum neural networks. *CoRR*, abs/2001.00550, 2020. URL http://arxiv.org/abs/2001.00550.

[61] Arthur Pesah, Marco Cerezo, Samson Wang, Tyler J. Volkoff, Andrew T. Sornborger, and Patrick J. Coles. Absence of barren plateaus in quantum convolutional neural networks. *CoRR*, abs/2011.02966, 2020. URL https://arxiv.org/abs/2011.02966.

[62] Michael Broughton, Guillaume Verdon, Trevor McCourt, Antonio J. Martinez, Jae Hyeon Yoo, Sergei V. Isakov, Philip Massey, Murphy Yuezhen Niu, Ramin Halavati, Evan Peters, Martin Leib, Andrea Skolik, Michael Streif, David Von Dollen, Jarrod R. McClean, Sergio Boixo, Dave Bacon, Alan K. Ho, Hartmut Neven, and Masoud Mohseni. Tensorflow quantum: A software framework for quantum machine learning. *CoRR*, abs/2003.02989, 2020. URL https://arxiv.org/abs/2003.02989.

[63] Hanrui Wang, Zhiding Liang, Jiaqi Gu, Zirui Li, Yongshan Ding, Weiwen Jiang, Yiyu Shi, David Z. Pan, Frederic T. Chong, and Song Han. Torchquantum case study for robust quantum circuits. In Tulika Mitra, Evangeline F. Y. Young, and Jinjun Xiong, editors, *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design, ICCAD 2022, San Diego, California, USA, 30 October 2022 - 3 November 2022*, pages 136:1–136:9. ACM, 2022. doi: 10.1145/3508352.3561118. URL https://doi.org/10.1145/3508352.3561118.

[64] Nam Nguyen and Kwang-Chen Chen. Quantum embedding search for quantum machine learning. *CoRR*, abs/2105.11853, 2021. URL https://arxiv.org/abs/2105.11853.

[65] Andrea Mari, Thomas R. Bromley, Josh A. Izaac, Maria Schuld, and Nathan Killoran. Transfer learning in hybrid classical-quantum neural networks. *Quantum*, 4:340, 2020. doi: 10.22331/Q-2020-10-09-340. URL https://doi.org/10.22331/q-2020-10-09-340.

[66] Maxwell P. Henderson, Samriddhi Shakya, Shashindra Pradhan, and Tristan Cook. Quanvolutional neural networks: Powering image recognition with quantum circuits. *CoRR*, abs/1904.04767, 2019. URL http://arxiv.org/abs/1904.04767.

[67] Amira Abbas, David Sutter, Christa Zoufal, Aurélien Lucchi, Alessio Figalli, and Stefan Woerner. The power of quantum neural networks. *Nat. Comput. Sci.*, 1(6):403–409, 2021. doi: 10.1038/S43588-021-00084-1. URL https://doi.org/10.1038/s43588-021-00084-1.

[68] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Leo Zhou. The quantum approximate optimization algorithm and the sherrington-kirkpatrick model at infinite size. *Quantum*, 6:759, 2022. doi: 10.22331/Q-2022-07-07-759. URL https://doi.org/10.22331/q-2022-07-07-759.

[69] Tingting Li, Ziming Zhao, Liqiang Lu, and Jianwei Yin. QuFrame: A novel encoding ensemble framework for quantum neural networks. In *2025 International Conference on Quantum Communications, Networking, and Computing (QCNC)*, pages 583–590. IEEE, 2025. doi: 10.1109/QCNC64685.2025.00097. URL https://doi.org/10.1109/QCNC64685.2025.00097.

[70] Tingting Li, Liqiang Lu, Ziming Zhao, Ziqi Tan, Siwei Tan, and Jianwei Yin. QuST: Optimizing quantum neural network against spatial and temporal noise biases. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 44(4):1434–1447, 2025. doi: 10.1109/TCAD.2024.3471949. URL https://doi.org/10.1109/TCAD.2024.3471949.

[71] Tingting Li, Ziming Zhao, and Jianwei Yin. Adaptive fidelity estimation for quantum programs with graph-guided noise awareness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 695–703, 2026. doi: 10.1609/aaai.v40i1.37035. URL https://doi.org/10.1609/aaai.v40i1.37035.

[72] Tingting Li, Ziming Zhao, and Jianwei Yin. Empowering quantum serverless circuit deployment optimization via graph contrastive learning and learning-to-rank co-designed approaches. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 9250–9258. International Joint Conferences on Artificial Intelligence Organization, 2025. doi: 10.24963/ijcai.2025/1028. URL https://doi.org/10.24963/ijcai.2025/1028.

[73] Ziming Zhao, Tingting Li, and Jianwei Yin. QuaMap: A multi-backend benchmark dataset for quantum circuit mapping and learning-based compiler evaluation. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1*, pages 2913–2923. ACM, 2026. doi: 10.1145/3770854.3785697. URL https://doi.org/10.1145/3770854.3785697.

## NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The contributions and the scope of the claims are stated in the abstract and Section 1, including the distinction between a physics-informed inductive bias and quantum hardware speedup. The construction is detailed in Section 3, and the empirical claims are reported in Tables 1–2 and Figures 4–5 in Section 4.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Section 6, in the Limitations paragraph, discusses dependence on the residual-mode dictionary and interaction graph, the use of quantum circuit simulation, hardware noise and connectivity, scalability, and the scope of suitable tasks.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification: The paper presents model definitions and a circuit construction in Sections 2–3, rather than formal theorems or theoretical guarantees requiring proofs.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [No]

   Justification: Sections 3 and 4.1 describe the architecture and evaluation protocol, but the manuscript does not provide all training hyperparameters, random seeds, and compute details needed to reproduce the reported experimental results. The implementation has not yet been released.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [No]

   Justification: Section 4.1 identifies the existing benchmark sources. The project repository at `https://github.com/fanqikong/NeurIPS2026-QRP` currently contains only a README stating that the code is coming soon; implementation files, configurations, and reproduction instructions have not yet been released.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [No]

   Justification: Section 4.1 gives the benchmarks, inputs, targets, baselines, and metrics, while Section 3.3 describes the circuit and objective. The manuscript does not specify the optimizer, learning rate, batch size, training epochs, or complete hyperparameter-selection procedure.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Tables 1–5 and Figures 4–5 report point estimates. The manuscript does not specify multi-seed variability, error bars, confidence intervals, or significance tests.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [No]

   Justification: Section 4.4 and Figure 5(b) report relative computational costs. The hardware, memory, storage, absolute run times, and total compute budget are not specified.

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The benchmark-based study and its input modalities are described in Section 4.1; Section 6 explains the quantum-simulation setting and deployment limitations. Within this stated research scope, the authors affirm compliance with the NeurIPS Code of Ethics to the best of their knowledge.

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [No]

    Justification: Sections 4 and 6 describe technical benefits and limitations, but the main text does not explicitly discuss positive and negative societal impacts or mitigation of potential harms from robot deployment.

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The work studies dynamics predictors on the existing benchmarks described in Section 4.1. It does not release high-risk generative models, scraped datasets, or models intended for direct unsupervised deployment in safety-critical settings.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [No]

    Justification: Section 4.1 credits the benchmark creators through citations, but the manuscript does not explicitly identify asset versions, license names, or applicable terms of use. Citations alone do not provide the complete licensing disclosure requested here.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [No]

    Justification: Section 3 describes the proposed method, but its implementation and release documentation have not yet been published. The repository linked in Item 5 currently contains only a code-coming-soon README.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The study uses existing benchmarks as described in Section 4.1 and does not conduct new crowdsourcing experiments or human-subject studies, so participant instructions and compensation are not applicable.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The benchmark-based study described in Section 4.1 does not involve new human-subject research or crowdsourcing, so IRB approval or equivalent participant review is not applicable.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: The core architecture and training procedure in Section 3 do not use LLMs as an important, original, or non-standard component. Use limited to writing, editing, or formatting does not affect the core methodology.