

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 ON THE COMPUTATIONAL LIMITS OF AI4S-RL : A UNIFIED ε - N ANALYSIS

Anonymous authors

Paper under double-blind review

ABSTRACT

Recent work increasingly adopts AI for Science (AI4S) models to replace expensive PDE solvers as simulation environments for reinforcement learning (RL), enabling faster training in complex physical control tasks. However, using approximate simulators introduces modeling errors that affect the learned policy. In this paper, we introduce a unified ε - N framework that quantifies the minimal computational cost $N^*(\varepsilon)$ required for an AI4S model to ensure that tabular RL can estimate the value function with unbiasedness, with probability at least $1 - \delta$. This characterization allows us to connect surrogate accuracy, grid resolution, and RL policy quality under a shared probabilistic language. We analyze how the discretization level K of AI4S and RL space governs both PDE surrogate error and RL lattice approximation error, and we employ spectral theory and Sobolev estimates to derive optimal grid strategies that minimize total cost while preserving learning fidelity. Our theory reveals that different systems – such as ODE- and PDE-governed environments – require different allocations of effort between physical simulation and RL optimization, which is consistent with the empirical results. Overall, our framework offers a principled foundation for designing efficient, scalable, and cost-aware AI4S-RL systems with provable learning guarantees.

1 INTRODUCTION

Reinforcement learning (RL) for PDE-constrained control promises transformative impact in domains such as fusion energy (Degrave et al., 2022) and climate modeling (Feng et al., 2025). Yet a fundamental obstacle remains: high-fidelity PDE simulations are far too costly to support the millions of interactions demanded by RL. A natural remedy is to replace PDEs with AI surrogates, but this introduces a critical and largely unexplored question: **how should the resolution of the surrogate be coordinated with the RL agent’s discretization so as to minimize total computational cost while preserving policy accuracy?**

Classical RL theory assumes *stochastic* transition noise that vanishes under repeated sampling (Sutton et al., 1998; Azar et al., 2017), and robust RL extends this to worst-case distributional uncertainty (Derman et al., 2021; Agarwal & Zhang, 2022). In contrast, AI4S surrogates incur *deterministic* errors stemming from spatial discretization, temporal integration, and boundary approximation. Such errors do not average out with more data and can systematically bias policy learning. In a cart-pole system, if the surrogate discretization does not adequately capture the upright equilibrium, the learned controller will consistently overshoot; regardless of how fine the RL action resolution is, policy accuracy is ultimately constrained by the surrogate’s resolution. Intuitively, minimizing regret under a fixed computational budget requires balancing the grid resolution of the AI4S surrogate with the decision resolution of the RL system.

To address this trade-off, we introduce a framework that explicitly models discretization-induced bias. By sampling multiple trajectories from perturbed initial conditions within observation uncertainty bounds, we recast deterministic surrogate errors as a statistical inference problem. This structured sampling bridges numerical analysis in AI4S with RL learning theory, enabling principled resolution coordination between surrogates and agents. We focus on tabular RL to establish a necessary condition for learnability, providing a resolution lower bound that guarantees feasibility and guides hyperparameter design for deep RL in AI4S systems.

Contributions. We introduce the ε - N framework that quantifies the minimal number of surrogate calls N required to achieve ε -accurate value estimation with confidence $1 - \delta$, assuming optimal resolution allocation between AI4S simulation and RL exploration. Our analysis reveals that this optimization problem exhibits surprising complexity: optimal resolution ratios follow non-obvious fractional scaling laws (e.g., $\Delta t_r = (\Delta t_w)^{1/3}$ for tokamak control, where Δt_r is the discrete timestep in the RL domain and Δt_w is that in the AI4S/World domain), have exponential sensitivity to system dynamics, and vary significantly across different physical systems. We provide:

- Resolution coupling theory:** Our ρ - K analysis shows that the forward prediction error ρ scales as $\mathcal{O}(1/H_r + 1/K^d)$ in the high-resolution regime, where H_r is the RL's temporal resolution (inverse decision frequency), K is the ratio of physical grid cells to RL states, and d is the spatial dimension of PDE system. This result quantifies how the resolutions of both the AI4S surrogate and the RL agent jointly determine the learning quality.
- Closed-form optimal allocations:** We derive system-specific expressions for the optimal resolution ratio K^* that minimize computational cost. For instance, $K^* = (7/4 \cdot \exp(1/\lambda_1))^{1/3}$ for tokamak plasma control and $K^* = 2 \cdot \exp(1/\lambda_1)$ for cart-pole systems, where λ_1 denotes the dominant modal growth rate of the PDE dynamics.
- Empirical validation:** Experiments confirm that both tabular RL and DQN exhibit the predicted non-monotonic sensitivity to resolution, with off-optimal parameter choices requiring computational cost scaling as approximately $N^{1.6}$ relative to the optimal N to reach the same accuracy. Moreover, there exists an optimal grid ratio K and an optimal balance between H_r and H_w , and these values are found to be relatively consistent across different learning algorithms.

Related Literature. Our work lies at the intersection of reinforcement learning (RL) theory and scientific computing. On the RL side, the tabular literature has established Probably Approximately Correct (PAC) bounds that characterize sample complexity (Sutton et al., 1998; Azar et al., 2017; Dann et al., 2017). Extensions to function approximation further reveal how this complexity depends on structural properties such as Bellman rank and eluder dimension (Jin et al., 2020; Wang et al., 2020; Jin et al., 2021). A key limitation of these analyses is their reliance on stochastic transition dynamics, where errors are assumed to average out. This assumption breaks down in scientific computing surrogates, where deterministic discretization errors introduce systematic, resolution-dependent biases. While robust RL frameworks (Derman et al., 2021; Agarwal & Zhang, 2022; Shi et al., 2023; Kwon et al., 2021) address model uncertainty through worst-case optimization, they primarily target irreducible stochastic noise. In contrast, our focus is on structured numerical biases, which we exploit to derive principled resolution–sample trade-offs.

On the scientific computing side, recent advances in neural solvers have accelerated PDE simulations by orders of magnitude. Neural operators (Li et al., 2020; Lu et al., 2021; Kovachki et al., 2023) and physics-informed neural networks (Raissi et al., 2019; Karniadakis et al., 2021) have enabled breakthroughs in weather prediction (Bi et al., 2023), turbulence modeling (Kochkov et al., 2021), and plasma control (Degrave et al., 2022). Although operator learning admits error bounds (Kovachki et al., 2023), prior work rarely investigates how surrogate resolution affects downstream RL performance. Classical PDE control methods rely on adjoints and Pontryagin's principle (Pontryagin, 2018), and recent applications of RL to PDEs (Farahmand et al., 2017; Han et al., 2018) typically assume access to true dynamics or ignore surrogate errors.

2 PRELIMINARIES

Notation. We study a hybrid system where an RL agent interacts with an AI4S surrogate approximating a PDE-governed environment. A comprehensive notation table is provided in Appendix A with key notations:

- Subscripts r and w denote RL and physical world parameters respectively.
- $K = \Delta x_r / \Delta x_w$ is the ratio of spatial resolutions between the RL and AI4S domains, where Δx_r is the spatial grid size in the RL domain and Δx_w is that in the AI4S domain.
- H_r and H_w denote the temporal discretization (time-step size) of the RL agent and the AI4S surrogate, respectively.

108 • ε, δ are the target accuracy and confidence levels.
 109

110 **AI4S Surrogate Errors.** AI4S models approximate PDE solvers through neural operators. For a
 111 PDE solution operator $\mathcal{G} : f \mapsto u$ approximated by surrogate \mathcal{G}_θ , the error decomposes as:
 112

$$113 \quad \|\mathcal{G}(f) - \mathcal{G}_\theta(f)\|_{L^2} \leq C_1 H^s + C_2 X^d \quad (2.1)$$

114 where H is temporal discretization, X is spatial resolution ratio, and s, d are parameters determined
 115 by the dimension and property of the underlying PDE solver.
 116

117 **Probably Approximately Correct (PAC) Learning in Tabular RL.** We define an episodic Markov
 118 Decision Process (MDP) as $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, R, H)$, where \mathcal{S} is a finite state space with $|\mathcal{S}| = S$, \mathcal{A}
 119 is a finite action space with $|\mathcal{A}| = A$, P is an unknown transition kernel, $R : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the
 120 reward function, and H is the episode horizon. For clarity, we omit the subscript r here; in later
 121 sections, we will use r to distinguish RL parameters from those of the underlying physical system.
 122

123 We call that an algorithm is PAC-MDP if for any $\varepsilon, \delta \in (0, 1)$, with probability at least $1 - \delta$, it outputs
 124 a policy $\hat{\pi}$ satisfying $V_1^*(s_0) - V_1^{\hat{\pi}}(s_0) \leq \varepsilon$ after at most $\text{poly}(S, A, H, 1/\varepsilon, \log(1/\delta))$ episodes.
 125 A classical result in tabular RL settings is that the UCB-VI algorithm Azar et al. (2017) achieves
 126 PAC-MDP guarantee with sample complexity

$$127 \quad N_{\text{episodes}} = \tilde{O} \left(\frac{SAH^3}{\varepsilon^2} \log \frac{1}{\delta} \right), \text{ where } \tilde{O} \text{ hides logarithmic factors.}$$

129 3 EMPIRICAL MOTIVATION: RESOLUTION COORDINATION IN AI4S-RL

131 Before developing our theoretical framework, we first empirically demonstrate that an optimal
 132 discretization scale exists in AI4S-RL systems. This balance between surrogate and agent resolutions
 133 follows a scaling relation across grids.
 134

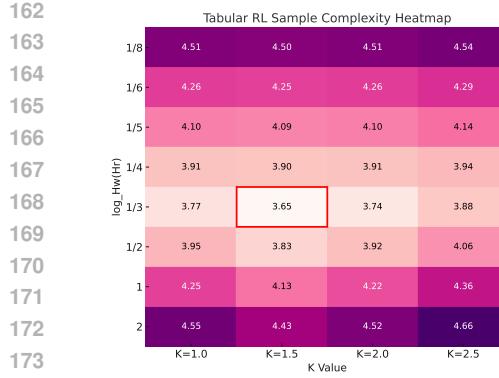
135 We utilize the simulated Cart-Pole environment, which is tractable for exhaustive search yet represen-
 136 tative of continuous dynamical systems. We establish two distinct experimental settings to bridge
 137 theory and practice. First, we employ Tabular Value Iteration on the exact physics engine to serve as
 138 a theoretical oracle. Second, to strictly emulate the AI4S workflow where learned dynamics replace
 139 expensive solvers, we train Deep Q-Networks (DQN) within a surrogate environment governed by a
 140 pre-trained neural network. Across both settings, we vary the temporal resolution ratio $\log_{H_w} H_r$ and
 141 the spatial resolution ratio K , measuring the sample complexity required to achieve a value function
 142 error below 1%. Full experimental details are provided in Appendix C.5.3.

143 Figures 1 and 2 show that sample complexity does not improve monotonically with resolution.
 144 Instead, both methods exhibit similar optima: tabular RL at $(K = 1.5, \log_{H_w} H_r = 1/3)$, and
 145 Q-learning at $(K = 2.0, \log_{H_w} H_r = 1/2)$. Off-optimal parameter choices require computational
 146 cost scaling as approximately $N^{1.6}$ relative to the optimal N in order to reach the same accuracy.

147 These findings highlight that conventional hyperparameter search strategies (e.g., grid or binary
 148 search) are ineffective in this setting: resolution parameters span multiple orders of magnitude,
 149 making brute-force exploration prohibitively expensive in high-dimensional PDE systems. This
 150 motivates the development of a theoretical framework that predicts optimal resolution trade-offs from
 151 system properties, rather than relying on empirical tuning.
 152

153 4 ERROR PROPAGATION AND RESOLUTION COORDINATION IN AI4S-RL 154 SYSTEMS

156 In this section, we develop a theoretical framework that predicts optimal resolution trade-offs from
 157 system properties. In Section 4.1, we show that the discretization of the AI4S surrogate itself
 158 determines whether the RL agent can converge to PAC guarantees. In Section 4.2, we analyze how
 159 the interaction between AI4S and RL discretization affects the convergence rate of the state transition
 160 matrix, using a magnetohydrodynamic (MHD) tokamak as a representative example. Building on
 161 these error models, Section 4.3 investigates the trade-off between the computational costs of surrogate
 construction and RL training. We derive the optimal resource allocation strategy—specifically the



174
175
176 Figure 1: Sample complexity heatmap for the
177 Cart-Pole system under tabular RL, showing non-
178 monotonic dependence on resolution parameters.

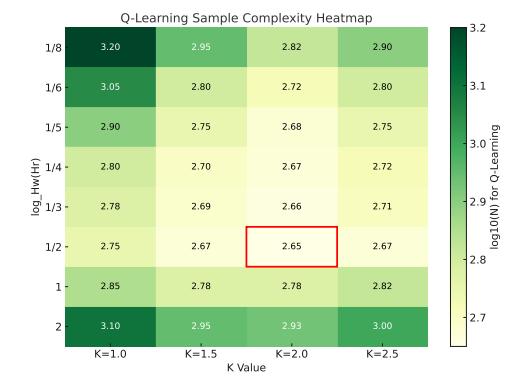


Figure 2: Sample complexity heatmap for the Cart-Pole system under Q-Learning, revealing algorithm-specific optimal configurations.

discretization resolutions—that minimizes the total cost required to achieve a PAC guarantee. Finally, in Section 4.4, we compute the optimal discretization and computational cost in four real-world systems to validate the framework.

4.1 SPECTRAL ANALYSIS OF ERROR AMPLIFICATION IN PDE-GOVERNED AI4S SYSTEMS

In AI4S systems governed by physical laws (e.g., PDEs), predictions are obtained by evolving the current state y_0 through a deterministic solver f under action a , yielding $y_1 = f(y_0, a)$. However, y_0 is only observed up to grid size Δy , so the predicted next state lies in the range:

$$[y_1^{\min}, y_1^{\max}] := \left[\min_{\|\Delta y_0\| \leq \Delta y} f(y_0 + \Delta y_0, a), \max_{\|\Delta y_0\| \leq \Delta y} f(y_0 + \Delta y_0, a) \right], \quad (4.1)$$

where Δy_0 is the measurement uncertainty in practice.

If the interval $[y_1^{\min}, y_1^{\max}]$ spans more than two full grid cells of PDE systems, the next state cannot be uniquely identified, and PAC guarantees are unattainable. In contrast, if the interval covers only a single full grid cell (with any additional coverage being partial), repeated sampling of the initial state allows us to estimate the next PDE state with arbitrarily high confidence. The convergence rate of this estimation is summarized in Theorem 1.

Theorem 1 (Sample Complexity for δ -Confidence Classification). *Consider repeated forward predictions from perturbed initial states, where the predicted frequency is the empirical probability of the next state falling into a given grid cell. Let p denote the predicted frequency with which the true next state falls into the correct grid cell, and let $q = p_{\max}^{(j)}$ denote the maximum predicted frequency among all competing cells. To resolve the correct cell with confidence at least $1 - \delta$, the number of forward predictions required is bounded by $n = \mathcal{O} \left(\frac{\log(1/\delta)}{\min_j (\Delta y^{(j)} - p_{\max}^{(j)})^2} \right)$*

Proof in Appendix C.1.1.

Remark 1 (Time-Step Constraint for State Separability). *Consider the solution operator $f_t(y_0)$ of a nonlinear PDE with Fréchet derivative $\mathcal{D}f_t(y_0)$. For observation perturbations η with $\|\eta\| \leq \Delta y$, the perturbation propagates as $\mathcal{D}f_t(y_0)[\eta] = \sum_k \hat{\eta}_k \gamma_k(t) \psi_k(t)$ where $\gamma_k(t) > 0$ are modal gain factors. Classification-cell separability requires $\sup_k \gamma_k(t) < 1$, implying a time-step constraint $\Delta t < \inf_k \gamma_k^{-1}(1)$.*

Proof in Appendix C.1.2.

From Theorem 1, when the two largest cell frequencies $\Delta y^{(j)}$ and $p_{\max}^{(j)}$ are close, where $(\Delta y^{(j)} - p_{\max}^{(j)})^2 \rightarrow 0$ and $n = \mathcal{O} \left(\frac{\log(1/\delta)}{(\Delta y^{(j)} - p_{\max}^{(j)})^2} \right) \rightarrow \infty$, so PAC is unattainable. Remark 1 shows that

smaller Δt enlarges the frequency gap ($\Delta y^{(j)} - p_{\max}^{(j)}$), while if $\Delta t > 1/\lambda_1$, perturbations exceed Δy and cell frequencies equalize, destroying state separability, where $\lambda_1 := \max_k\{\text{Re}(\lambda_k)\}$ denotes the dominant modal growth rate (i.e., the largest real part of the eigenvalues) of the system's Fréchet linearization, representing the intrinsic instability. Hence the constraint $\Delta t \lesssim 1/\lambda_1$ is necessary for PAC guarantees. We note that defining λ_1 as the global maximum is primarily for theoretical parsimony; utilizing a step-specific local growth rate in practice would not fundamentally alter the complexity analysis.

4.2 ERROR COUPLING BETWEEN RL AND PDE SPACES IN AI4S SYSTEMS

In Section 4.1, we showed that the intrinsic dynamics and temporal discretization of the physical system determine whether PAC guarantees are achievable, and computed the sample complexity to reach PAC when possible.

In this section, we extend the analysis to mismatched discretizations between the RL and PDE spaces. Each RL action a_r must be lifted to the PDE space for forward evolution, and the resulting physical state projected back for RL. These projections introduce errors, which may be amplified near PDE boundaries by nonlinear operators. Since RL actions often operate directly on boundaries (e.g., controlling droplet interfaces in tokamaks to avoid wall contact, or shaping fluid boundaries around an aircraft wing to increase lift), boundary accuracy is crucial. In MHD systems with droplet interfaces, the trace theorem implies that poor boundary approximation further enlarges these errors. To clarify the relationship between the two spaces, we define two discretization structures:

- RL space: (H_r, S_r, A_r) , representing the agent's temporal discretization, state discretization, and action resolution.
- PDE space: (H_w, S_w, A_w) , representing the physical system's temporal discretization, state resolution, and control parameter resolution.

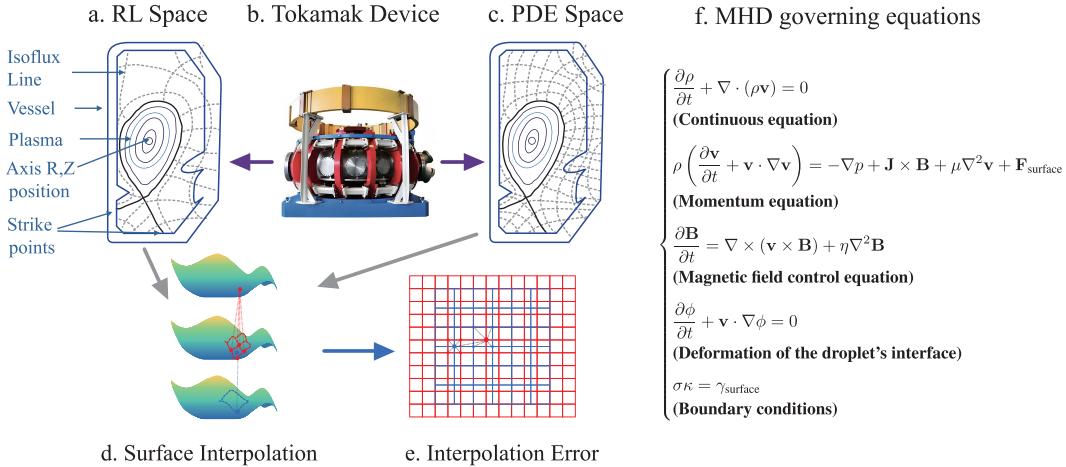


Figure 3: Illustration of Error Propagation in a Tokamak Machine RL controlling System.

MHD System As a representative case, in this section we focus on a magnetohydrodynamic (MHD) tokamak, which enables real-time magnetic control to sustain plasma droplet confinement without wall contact (Degrave et al., 2022; Ding et al., 2024). We assume full observation of the MHD fluid and magnetic field, and conduct numerical validation based on realistic tokamak configurations.

Figure 3 illustrates the governing PDEs, interface tracking methods, and control structure involved in the RL-based regulation of the droplet. The reward signal for reinforcement learning is defined as the minimum distance between the plasma droplet and the tokamak boundary, and the control actions correspond to adjustments of the magnetic field via coil currents. These equations define the error sources and discretization structure that form the basis for our multi-resolution analysis.

We assume continuous control over the magnetic field via time-dependent coil currents $\mathbf{I}_{\text{coil}}(t)$, enabling real-time adjustments to maintain droplet suspension without wall contact. Reinforcement learning algorithms adapt these control inputs based on observations of the droplet boundary and plasma dynamics. Key physical variables—including the magnetic field \mathbf{B} , current density \mathbf{J} , density ρ , velocity field \mathbf{v} , and interface function ϕ —are continuously monitored and regulated to optimize positioning and stability.

Interface Error Propagation via Control-Boundary Coupling In tokamak control, RL agents manipulate boundary magnetic coils to suspend plasma droplets away from walls. This indirect control—propagating from boundaries through multiple physical fields to the interface—introduces compounding errors. We first analyze single-equation timestep errors, then demonstrate their coupling across the full system.

The control-to-interface coupling follows a hierarchical propagation path:

$$\text{RL action } a_r \xrightarrow{\text{projection}} \Delta a_w \xrightarrow{\text{boundary}} \mathbf{B}|_{\partial\Omega} \xrightarrow{\text{curl}} \mathbf{J} \xrightarrow{\text{Lorentz}} \mathbf{v} \xrightarrow{\text{advection}} \phi \quad (4.2)$$

where each arrow represents a potential error amplification point. The magnetic field \mathbf{B} controlled at the boundary generates current density $\mathbf{J} = \nabla \times \mathbf{B}$, which produces Lorentz forces that drive the velocity field \mathbf{v} , ultimately advecting the interface level set ϕ .

At each stage, discretization errors accumulate. The RL action a_r must be mapped to physical control parameters a_w , introducing projection error $\mathcal{O}(\Delta a_r - \Delta a_w)$. Magnetic field boundary conditions suffer from reduced regularity, yielding $\|\mathbf{B} - \mathbf{B}_h\|_{L^2(\partial\Omega)} = \mathcal{O}(\|\Delta \mathbf{x}_{p,\text{bd}}\|^{1/2})$ by the trace theorem. These control errors then propagate nonlinearly into the velocity field as $\Delta \mathbf{v} = \mathcal{O}(\Delta a_w^{1/2})$, affecting the entire flow dynamics.

The droplet interface evolves according to the level set equation $\frac{\partial \phi}{\partial t} + \mathbf{v} \cdot \nabla \phi = 0$. Under forward Euler time discretization and central spatial differences, the local truncation error per PDE timestep k becomes $\tau^{(k)} = \mathcal{O}(\Delta t_w) + \mathcal{O}(\|\Delta \mathbf{x}_{p,\text{int}}\|)$.

However, incorporating the velocity perturbations from boundary control and the reduced regularity at interfaces, the complete per-step interface error is:

$$\Delta_{\phi}^{(k)} = \underbrace{\mathcal{O}(\Delta t_w)}_{\text{temporal}} + \underbrace{\mathcal{O}(\|\Delta \mathbf{x}_{p,\text{int}}\|)}_{\text{spatial interior}} + \underbrace{\mathcal{O}(\Delta a_w^{1/2})}_{\text{control}} + \underbrace{\mathcal{O}(\|\Delta \mathbf{x}_{p,\text{bd}}\|^{1/2})}_{\text{boundary}} + \underbrace{\mathcal{O}(\Delta a_w^{1/2} \cdot \|\Delta \mathbf{x}_{p,\text{bd}}\|^{1/2})}_{\text{coupled}} \quad (4.3)$$

In the following paragraph, we will couple this error with contributions from other governing equations to enable a comprehensive analysis.

Refined Total Error Decomposition with Time-Scale Separation By extending the interface error structure discussed previously, we generalize the analysis to a broader class of PDE-based surrogate environments (details in Supplementary Information). Over a single RL step, the total prediction error can be decomposed as:

$$\begin{aligned} \Delta_{\text{total}} = & \underbrace{C_1 \|\Delta \mathbf{x}_{r,\text{int}}\|}_{\text{RL (interior space)}} + \underbrace{C_2 \|\Delta \mathbf{x}_{r,\text{bd}}\|^{1/2}}_{\text{RL (boundary space)}} + \underbrace{C_3 \Delta a_r^{1/2}}_{\text{RL (action space)}} + \underbrace{C_4 \Delta t_r}_{\text{RL (time)}} \\ & + \underbrace{C_5 \frac{\Delta t_r}{\Delta t_w} \left(\|\Delta \mathbf{x}_{p,\text{int}}\| + \Delta a_w^{1/2} + \|\Delta \mathbf{x}_{p,\text{bd}}\|^{1/2} + \Delta a_w^{1/2} \cdot \|\Delta \mathbf{x}_{p,\text{bd}}\|^{1/2} \right)}_{\text{PDE surrogate error (space + action + boundary), amplified by time scale separation}}. \end{aligned} \quad (4.4)$$

To ensure that the RL transition kernel remains distinguishable under observation uncertainty Δy , we impose the condition $\Delta_{\text{total}} = \mathcal{O}(\Delta y)$. This leads to the following resolution matching constraints between RL and AI4S components:

$$\begin{aligned} \|\Delta \mathbf{x}_{r,\text{int}}\| & \sim \|\Delta \mathbf{x}_{p,\text{int}}\| \sim \Delta a_r^{1/2} \sim \Delta a_w^{1/2} \sim \Delta t_r \sim \Delta y, \\ \|\Delta \mathbf{x}_{r,\text{bd}}\| & \sim \|\Delta \mathbf{x}_{p,\text{bd}}\| \sim \Delta y^2, \quad \Delta t_w \ll \Delta y. \end{aligned} \quad (4.5)$$

This decomposition makes explicit how discretization scales in the RL and AI4S pipelines must be jointly selected to ensure unbiased value estimation. In particular, the term $\Delta t_r/\Delta t_w$ highlights how

324 fine-scale physical errors can be amplified over coarser RL horizons, motivating resolution-aware
 325 control design.

326 In Section 4.1, we analyzed how to approximate the RL transition kernel under observation noise Δy .
 327 Here, we incorporate numerical errors from the AI4S prediction and back. Concretely, let $p(c^*)$ denote the possibly that the true
 328 next physical state given one observed state falls into the correct grid cell c^* after (i) projecting the
 329 observed PDE state into the RL grid, (ii) selecting the truncated optimal RL action, and (iii) mapping
 330 this action back to the PDE. Then $\rho := 1 - p(c^*)$, which quantifies the probability that one RL-PDE
 331 interaction fails to preserve grid-level consistency.

332 From Eq. 4.4 and 4.5, total numerical error $\Delta_{\text{total}} \sim \frac{C_1}{K^3} \Delta y \sim \frac{C_1}{K^3} \Delta t_r$, and letting the RL temporal
 333 resolution be $H_r = 1/\Delta t_r$ (we also use this further on as a unit for calculating total computational
 334 cost), the total error rate is:

$$335 \rho = 1 - \frac{\Delta y}{\lambda_1 \Delta y / H_r + \Delta y + \Delta_{\text{total}}} = 1 - \frac{1}{\lambda_1 / H_r + 1 + C_1 / K^3}. \quad (4.6)$$

336 Hence, in the high-resolution limit (when H_r and K^3 are large) $\rho = \mathcal{O}\left(\frac{1}{H_r} + \frac{1}{K^3}\right)$, showing that
 337 finer RL temporal resolution and AI4S spatial refinement both reduce the forward projection error
 338 rate at inverse polynomial rates.

339 From more general analysis on PDE scales, we could get:

340 **Theorem 2** (ρ - K Analysis for d -Dimensional Systems). *For a d -dimensional PDE system, the
 341 forward projection error rate is: $\rho = 1 - \frac{1}{\lambda_1 / H_r + 1 + C_1 / K^d}$, where the numerical error scales as
 342 $\Delta_{\text{total}} \sim C_1 K^{-d} \Delta y$ from spatial discretization. In the high-resolution limit:*

$$343 \rho = \mathcal{O}\left(\frac{1}{H_r} + \frac{1}{K^d}\right).$$

344 *Proof in Appendix C.1.4.*

345 4.3 OPTIMAL COMPUTATIONAL COST ALLOCATION BETWEEN RL AND AI4S

346 To translate these PAC-possible AI4S-RL analysis into practical system design guidance, this section
 347 establishes a computational resource allocation framework for AI4S-RL systems. Since AI4S models
 348 are generally trained under fixed computational budgets, we aim to derive how their discretization
 349 parameters—denoted as H_w, S_w, A_w —can be aligned with the RL-side parameters H_r, S_r, A_r to
 350 ensure computational and statistical consistency.

351 On the physical side, computational cost scales as $H_w \cdot S_w \cdot A_w$. For tabular RL, classical UCB-
 352 VI methods require a sample complexity of $\mathcal{O}(H_r^4 S_r A_r)$ for value function convergence without
 353 any assumptions on transition smoothness. Based on our transition identifiability analysis, an
 354 improved bound of $\mathcal{O}(H_r S_r A_r \cdot \frac{\log(1/\delta)}{\min_j (\Delta y^{(j)} - p_{\max}^{(j)})^2})$ is achievable when the transition matrix is
 355 known up to statistical confidence δ , reducing the RL problem to dynamic programming. From our
 356 ρ - K analysis, we approximate $p_{\max}^{(j)} \lesssim \rho \Delta y \sim \frac{1}{H_r} \Delta y$, leading to a refined sample complexity of
 357 $\mathcal{O}(H_r^3 S_r A_r \cdot \log(1/\delta))$, which we adopt in the remainder of this section.

358 We then align RL and AI4S computational costs via the relation $H_w S_w A_w = H_r^3 S_r A_r$. From the
 359 interface and grid-based error analysis in Section 4.2 for MHD system, and the temporal resolution
 360 bound in Section 4.1, we discretize the physical space in three spatial dimensions (x, y, z) , where:
 361 $A_r = \mathcal{O}(A_w) = \mathcal{O}(x_w^2)$, $S_r = \mathcal{O}(S_w) = \mathcal{O}(x_w^6)$, $H_r = \mathcal{O}(H_w^{1/3})$, $H_w = \mathcal{O}(x_w) \lesssim \frac{1}{\lambda_1}$. This
 362 constraint also respects CFL stability conditions for fluid systems.

363 By substituting these scaling relations into our cost-balancing equation, we can express the total
 364 computational cost as a function of the resolution ratio K , leading to the following theorem.

365 **Theorem 3** (Optimal Resolution with System-Dependent Scaling). *For a physical system subject
 366 to projection error with state space scaling $S_r \sim K^\alpha$ and action space scaling $A_r \sim K^\beta$, the com-
 367 putational balance condition $H_r^3 S_r A_r \sim H_w S_w A_w$ ensures equivalence between the RL surrogate
 368 resolution and the underlying PDE world resolution. Minimizing the overall computational costs,*

378
379
380
381

$$Cost(K) = H_r^3 K^{\alpha+\beta} \cdot \left(\frac{\log(1/\delta)}{\varepsilon^2} \right) \cdot \left(\frac{1}{1 - \frac{1}{H_r} - \frac{1}{K^d}} \right)^2 \quad (4.7)$$

382 yields optimal resolution ratio between RL and AI4S space:

383
384
385
386
387

$$K^* = \left(\frac{\alpha + \beta + 2d}{(\alpha + \beta)(1 - H_r^{-1})} \right)^{1/d} \approx \left(\frac{\alpha + \beta + 2d}{\alpha + \beta} \right)^{1/d} \cdot \exp \left(\frac{1}{d\lambda_1} \right) \quad (4.8)$$

when $H_r \gtrsim \lambda_1 \gg 1$.388
389
390
391
392
393
394
395

Proof in Appendix C.1.5.

Remark 2 (Impact of Actuation Topology). The scaling exponents derive from the Trace Theorem based on the control location. **Boundary-actuated systems** (e.g., tokamak) require quadratic refinement due to boundary regularity loss, yielding $\alpha = 2d$ and $\beta = 2d_a$. Conversely, **interior-actuated systems** (e.g., heat sequencing) follow standard volumetric scaling with $\alpha = d$ and $\beta = d_a$, where d denotes the physical spatial dimension and d_a the action dimension. This distinction dictates the optimal resolution K^* and the cost disparities shown in Table 1.

396
397
398
399
400
401
402

4.4 THEORETICAL ANALYSIS ACROSS PHYSICAL SYSTEMS

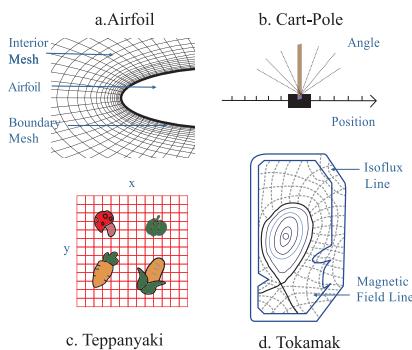
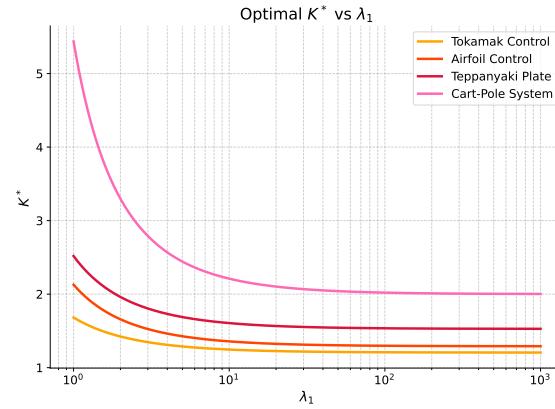
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417

Figure 4: State space discretization for four AI4S-RL systems.

Figure 5: Log-scale plot of K^* vs. modal growth λ_1 .421
422
423
424
425
426

We adopted an ε - N perspective: to ensure value function estimation within error ε and confidence $1 - \delta$, how many AI4S-RL transitions N are required, given surrogate discretization and physical uncertainty? Based on our analysis, we can now systematically achieve optimal resolution allocation under (ε, δ) accuracy constraints through a unified framework that integrates error normalization, identifiability analysis, and closed-form cost expressions—thereby minimizing the total computational cost of AI4S-RL systems while ensuring reliable policy estimation.

427
428
429
430
431

Table 1 summarizes scaling relations and optimal resolution ratios. Systems with strong boundary observability (e.g., tokamak, airfoil) require quadratic scaling in state-action discretization relative to mesh resolution, while low-dimensional systems tolerate aggressive upsampling without cost explosion. Figure 5 shows how the optimal K^* saturates with increasing λ_1 , especially in high-dimensional PDEs. In contrast, ODE-based systems benefit more directly from higher surrogate resolution.

Our theoretical framework provides insights into the empirical results from Section 3. For tabular RL, the observed optimum at $(K = 1.5, \log_{H_w} H_r = 1/3)$ requires $10^{3.65} \approx 4500$ samples, while the worst configuration at $(K = 2.5, \log_{H_w} H_r = 2)$ needs $10^{4.66} \approx 46000$ samples—a 10-fold increase. This translates to computational cost scaling approximately as $N^{1.6}$ when deviating from optimal parameters, consistent with our claim in the abstract. Q-learning achieves its minimum at $(K = 2.0, \log_{H_w} H_r = 1/2)$ with only $10^{2.65} \approx 450$ samples, demonstrating that function approximation can effectively exploit state similarity to reduce sample complexity by an order of magnitude.

Table 1: Optimal Resolution Matching and Computation Cost Across Systems under ε error of value function estimation and $1 - \delta$ confidence level

System	Resolution Scaling	Computational Cost	Optimal K^*
Tokamak Control	$A_r = \mathcal{O}(A_w) = \mathcal{O}(x_w^2)$ $S_r = \mathcal{O}(S_w) = \mathcal{O}(x_w^6)$ $H_r = H_w^{1/3} \cdot \mathcal{O}(x_w) < 1/\lambda_1$	$\frac{H_r^9 K^8}{\varepsilon^2} \cdot \frac{\log(1/\delta)}{(1 - \frac{1}{H_r} - \frac{1}{K^3})^2}$	$\left(\frac{7}{4}\right)^{1/3} \cdot \exp\left(\frac{1}{3\lambda_1}\right)$
Airfoil Control	$A_r = \mathcal{O}(A_w) = \mathcal{O}(x_w^2)$ $S_r = \mathcal{O}(S_w) = \mathcal{O}(x_w^4)$ $H_r = H_w^{1/3} \cdot \mathcal{O}(x_w) < 1/\lambda_1$	$\frac{H_r^7 K^6}{\varepsilon^2} \cdot \frac{\log(1/\delta)}{(1 - \frac{1}{H_r} - \frac{1}{K^2})^2}$	$\left(\frac{5}{3}\right)^{1/2} \cdot \exp\left(\frac{1}{2\lambda_1}\right)$
Teppanyaki Plate	$A_r = \mathcal{O}(A_w) = \mathcal{O}(x_w)$ $S_r = \mathcal{O}(S_w) = \mathcal{O}(x_w^2)$ $H_r = H_w^{1/3} \cdot \mathcal{O}(x_w) < 1/\lambda_1$	$\frac{H_r^4 K^3}{\varepsilon^2} \cdot \frac{\log(1/\delta)}{(1 - \frac{1}{H_r} - \frac{1}{K^2})^2}$	$\left(\frac{7}{3}\right)^{1/2} \cdot \exp\left(\frac{1}{2\lambda_1}\right)$
Cart-Pole System	$A_r = \mathcal{O}(A_w) = \mathcal{O}(x_w)$ $S_r = \mathcal{O}(S_w) = \mathcal{O}(x_w)$ $H_r = H_w^{1/3} \cdot \mathcal{O}(x_w) < 1/\lambda_1$	$\frac{H_r^2 K^2}{\varepsilon^2} \cdot \frac{\log(1/\delta)}{(1 - \frac{1}{H_r} - \frac{1}{K})^2}$	$2 \cdot \exp\left(\frac{1}{\lambda_1}\right)$

These results suggest that systems with high spectral amplification yield diminishing returns from AI4S refinement unless RL granularity is co-optimized. Our ε - N framework thus provides a principled tool for balancing simulation fidelity and RL efficiency under finite compute budgets.

4.5 EMPIRICAL VALIDATION ON PDE-GOVERNED CONTROL: TEPPANYAKI HEAT SEQUENCING

While the Cart-Pole experiments in Section 3 successfully demonstrate resolution trade-offs, they rely on a low-dimensional ODE system with learned dynamics. To rigorously validate our theoretical predictions in high-dimensional PDE environments with continuous state-action spaces, we conduct a comprehensive study on a two-dimensional control problem: multi-item heat sequencing on a teppanyaki cooking surface. This environment is governed by the canonical two-dimensional heat diffusion equation $\partial T / \partial t = \alpha \nabla^2 T + Q(\mathbf{x}, t)$, representing a broad class of parabolic PDEs encountered in AI4S applications (thermal management, plasma diffusion, etc.) (details in Appendix E).

The control task requires simultaneously cooking three food items at distinct temperatures using two heat sources, posing challenges for policy learning under discretization constraints. We train Proximal Policy Optimization (PPO) agents across a Cartesian product of spatial resolution ratios K and temporal resolution parameters $\log_{h_w}(h_r)$, yielding 25 experimental configurations as detailed in Figure 6. All experiments run on identical hardware (Intel Xeon Gold 6530, NVIDIA RTX 4090). Complete environment specifications, PPO hyperparameters, and convergence criteria are provided in Appendix E.

Figure 6 illustrates the comprehensive experimental results on the Teppanyaki thermal control task. Panels (a) and (b) present the environment setup: the agent regulates distributed heat sources to manage non-stationary diffusion dynamics for precise multi-object temperature tracking. Panel (c) displays ε - N scaling curves for all 23 qualified configurations, where color denotes temporal resolution y and intensity indicates spatial resolution K . Panels (d), (e), and (f) show CPU, GPU, and total computational cost heatmaps across the (K, y) configuration space. Configurations at $K = 8.0$ (marked "R.I.") fail to converge, validating that excessive coarsening violates learnability conditions. The optimal configuration ($K^* = 6, \log_{h_w}(h_r) = 1/3$) achieves minimum total cost, confirming that Theorem 3 provides effective guidance for deep RL hyperparameter selection.

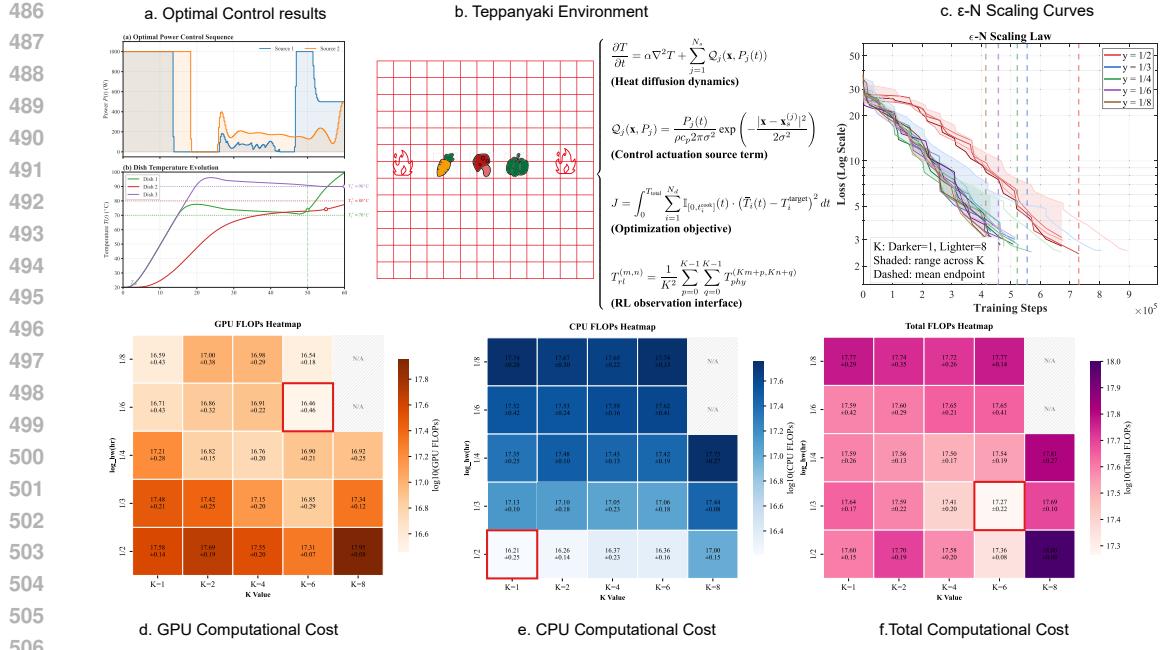


Figure 6: Experimental analysis of the Teppanyaki PDE control task. (a) The analytical optimal control solution serves as a performance oracle, illustrating the precise temperature tracking required. (b) Visualization of the simulation environment and the governing MHD/heat diffusion equations. (c-e) ϵ -N scaling curves (top) with CPU/GPU computational cost decomposition heatmaps (bottom), demonstrating the non-monotonic resolution-efficiency trade-off.

5 CONCLUSION AND LIMITATIONS

This work presents a unified theoretical framework for analyzing RL in surrogate-based physical environments governed by PDEs. By integrating spectral discretization theory with PAC sample complexity bounds from tabular RL, we establish how surrogate resolution, physical dynamics, and RL precision jointly influence the number of episodes required for accurate policy learning.

We propose the ϵ -N framework to characterize the minimal RL-AI4S interaction cost under fixed accuracy and confidence requirements. Through analysis of surrogate-induced error propagations, and optimal resolution trade-off, we derive system-specific cost scaling laws and identify spectral regimes where surrogate refinement provides diminishing returns. Empirical validation across four representative AI4S-RL systems confirms that spectral growth rate, boundary observability, and dimensionality all play critical roles in determining learning efficiency.

Several assumptions in our analysis warrant further investigation. First, our analysis focuses on tabular RL for theoretical tractability and to isolate physical discretization errors, establishing a necessary condition for learnability that applies to any algorithm. While our experiments with DQN and PPO confirm that these resolution trade-offs govern the primary performance trends even in deep RL, rigorously extending our mathematical bounds to function approximation settings remains an open direction. Second, we treat surrogate models as fixed and externally trained; adaptive refinement or active error correction within the RL loop could further improve performance. Finally, our framework focuses on deterministic discretization errors. While real-world neural surrogates also introduce approximation and generalization errors, our analysis establishes a theoretical baseline, isolating the resolution trade-off as a fundamental constraint.

Overall, this work provides a principled foundation for evaluating and co-designing AI4S surrogates and RL policies in resource-constrained scientific applications.

540 REFERENCES

542 Alekh Agarwal and Tong Zhang. Model-based rl with optimistic posterior sampling: Structural
543 conditions and sample complexity. *Advances in Neural Information Processing Systems*, 35:
544 35284–35297, 2022.

545 Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforce-
546 ment learning. In *International conference on machine learning*, pp. 263–272. PMLR, 2017.

548 Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. Accurate medium-
549 range global weather forecasting with 3d neural networks. *Nature*, 619(7970):533–538, 2023.

550 Christoph Dann, Tor Lattimore, and Emma Brunskill. Unifying pac and regret: Uniform pac bounds
551 for episodic reinforcement learning. *Advances in Neural Information Processing Systems*, 30,
552 2017.

554 Jonas Degrave, Federico Felici, Jonas Buchli, et al. Magnetic control of tokamak plasmas through
555 deep reinforcement learning. *Nature*, 602(7895):414–419, 2022.

556 Esther Derman, Matthieu Geist, and Shie Mannor. Twice regularized mdps and the equivalence
557 between robustness and regularization. *Advances in Neural Information Processing Systems*, 34:
558 22274–22287, 2021.

560 Shiping Ding, AM Garofalo, HQ Wang, DB Weisberg, ZY Li, X Jian, David Eldon, BS Victor,
561 Alessandro Marinoni, QM Hu, et al. A high-density and high-confinement tokamak plasma regime
562 for fusion energy. *Nature*, 629(8012):555–560, 2024.

563 Weinan E. *Principles of Multiscale Modeling*. Cambridge University Press, 2011.

564 Lawrence C. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*.
565 American Mathematical Society, 1998.

567 Amir-massoud Farahmand, Saleh Nabi, and Daniel N Nikovski. Deep reinforcement learning for
568 partial differential equation control. In *2017 American Control Conference (ACC)*, pp. 3120–3127.
569 IEEE, 2017.

571 Kairui Feng, Ning Lin, Robert E Kopp, Siyuan Xian, and Michael Oppenheimer. Reinforcement
572 learning–based adaptive strategies for climate change adaptation: An application for coastal flood
573 risk management. *Proceedings of the National Academy of Sciences*, 122(12):e2402826122, 2025.

574 Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations
575 using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.

577 Thomas Y. Hou. Numerical approximations to multiscale solutions in partial differential equations.
578 In J. F. Blowey, A. W. Craig, and T. Shardlow (eds.), *Frontiers in Numerical Analysis*, pp. 175–234.
579 Springer, 2003.

580 Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement
581 learning with linear function approximation. In *Conference on learning theory*, pp. 2137–2143.
582 PMLR, 2020.

583 Chi Jin, Qinghua Liu, and Sobhan Miryoosefi. Bellman eluder dimension: New rich classes of rl
584 problems, and sample-efficient algorithms. *Advances in neural information processing systems*,
585 34:13406–13418, 2021.

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

589 Dmitrii Kochkov, Jamie Smith, et al. Machine learning–accelerated computational fluid dynamics.
590 *Proceedings of the National Academy of Sciences*, 118(21):e2101784118, 2021.

592 Felix Koehler, Simon Niedermayr, Nils Thuerey, et al. Apebench: A benchmark for autoregressive
593 neural emulators of pdes. *Advances in Neural Information Processing Systems*, 37:120252–120310,
2024.

594 Nikola B. Kovachki, Lorenzo Rosasco, and Andrew M. Stuart. Neural operator learning: From theory
 595 to applications. *Foundations and Trends in Machine Learning*, 16(3):200–402, 2023.
 596

597 Jeongyeol Kwon, Yonathan Efroni, Constantine Caramanis, and Shie Mannor. RI for latent mdps:
 598 Regret guarantees and a lower bound. *Advances in Neural Information Processing Systems*, 34:
 599 24523–24534, 2021.

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

603 Jacques-Louis Lions. *Quelques méthodes de résolution des problèmes aux limites non linéaires*.
 604 Dunod; Gauthier-Villars, 1969.

605 Lu Lu, Pengzhan Jin, Guofei Pang, Zhen Zhang, and George Em Karniadakis. Learning nonlinear
 606 operators via deeponet based on the universal approximation theorem of operators. *Nature Machine
 607 Intelligence*, 3(3):218–229, 2021.

608 Andrew J. Majda and Andrea L. Bertozzi. *Vorticity and Incompressible Flow*, volume 27 of *Cam-
 609 bridge Texts in Applied Mathematics*. Cambridge University Press, 2002.

610 Savinay Nagendra, Nikhil Podila, Rashmi Ugarakhod, and Koshy George. Comparison of reinforce-
 611 ment learning algorithms applied to the cart-pole problem. In *2017 international conference on
 612 advances in computing, communications and informatics (ICACCI)*, pp. 26–32. IEEE, 2017.

613 Lev Semenovich Pontryagin. *Mathematical theory of optimal processes*. Routledge, 2018.

614 Koldo Portal-Porras, Unai Fernandez-Gamiz, Ekaitz Zulueta, Alejandro Ballesteros-Coll, and Asier
 615 Zulueta. Cnn-based flow control device modelling on aerodynamic airfoils. *Scientific Reports*, 12
 616 (1):8205, 2022.

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

620 Laixi Shi, Gen Li, Yuting Wei, Yuxin Chen, Matthieu Geist, and Yuejie Chi. The curious price of
 621 distributional robustness in reinforcement learning with a generative model. *Advances in Neural
 622 Information Processing Systems*, 36:79903–79917, 2023.

623 Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*, volume 1. MIT
 624 press Cambridge, 1998.

625 Ruosong Wang, Russ R Salakhutdinov, and Lin Yang. Reinforcement learning with general value
 626 function approximation: Provably efficient approach via bounded eluder dimension. *Advances in
 627 Neural Information Processing Systems*, 33:6123–6135, 2020.

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 **A TABLE OF NOTATIONS**
649650
651 The following table provides a reference for the key symbols and notations used throughout this
652 paper.
653

655 Symbol	655 Description
General Reinforcement Learning (RL) Parameters	
S, \mathcal{A}	The finite sets representing the state space and action space.
H	The horizon, representing the length of an episode.
$ \mathcal{S} , \mathcal{A} $	The size of the state and action spaces.
ε	The desired error tolerance for value function estimation.
δ	The confidence parameter.
$V^\pi(s)$	The value function of a policy π at state s .
V^*	The optimal value function.
$\pi^*, \hat{\pi}$	The optimal policy and the learned policy, respectively.
P, \hat{P}	The true and the empirically estimated state transition kernels.
N	The sample complexity.
N The overall computational cost is treated as a function of N in our ε - N analysis.	
AI4S-RL Framework Parameters	
S_r, A_r, H_r	The scale of the state space, action space, and horizon in the RL domain.
S_w, A_w, H_w	The scale of the state space, action space, and horizon in the AI4S domain.
$\Delta t_r, \Delta t_w$	The discrete timesteps in the RL and AI4S/World domains, respectively.
$\Delta x_r, \Delta x_w$	The spatial grid resolution in the RL and AI4S/PDE domains.
K	The Resolution Ratio, defined as $K = \Delta x_r / \Delta x_w$.
K^*	The optimal resolution ratio that minimizes computational cost.
$y_0, \Delta y_0$	An initial state and the measurement uncertainty in practice.
Δy	The grid cell size.
Δ_{total}	The total numerical error accumulated in one RL-AI4S interaction step.
ρ	The forward projection error rate in the AI4S-RL system.
PDE and Spectral Analysis Parameters	
λ_1	The dominant (largest real part) eigenvalue of the linearized PDE operator.
$\mathcal{L}(y_0)$	The Fréchet-linearized operator of the PDE system around a state y_0 .
ϕ_k	The k -th orthonormal basis function.
$\gamma_k(t)$	The modal gain factor for the k -th mode at time t .
$\psi_k(t)$	The propagated mode shape for the k -th mode at time t .
Mathematical Spaces and Operators	
$H^s(\Omega)$	The Sobolev space of order s over a domain Ω .
$L^p(\Omega)$	The Lebesgue space of p -integrable functions over a domain Ω .
∇, Δ	The gradient and Laplacian operators, respectively.
$\ \cdot\ _V$	The norm in a vector space V .
$\mathcal{D}f(y)[\eta]$	The Fréchet derivative of an operator f at point y applied to a perturbation η .

692
693
694 **B SAMPLE COMPLEXITY IN REINFORCEMENT LEARNING**
695696
697 In this section, we provide an explanation of sample complexity in reinforcement learning. In our
698 setting, the next-state information required by the reinforcement learning process is provided by a
699 PDE-based physical space. We assume that this physical space satisfies the property of probabilistic
700 determinability of the future state, meaning that, as long as repeated experiments are sufficiently
701 performed, the next state can be identified with probability at least one. We present a detailed
explanation of the probabilistic determinability of the future state.

702 B.1 DIRECTIONAL PERTURBATION IN A NONLINEAR DYNAMICAL SYSTEM
703

704 To understand uncertainty propagation in complex dynamical systems, consider a nonlinear system
705 governed by a PDE-based evolution operator $f(y_0, a)$, where y_0 is the initial state and a denotes
706 system parameters or external inputs. We focus on a single dimension (or component) of the system
707 state, say $y^{(i)}$, and analyze how small perturbations in the initial condition affect the predicted future
708 value $y_1^{(i)}$. Suppose that we do not observe y_0 exactly, but only a perturbed version $y_0 + \Delta y_0$, where
709 $\|\Delta y_0\| \leq \Delta y$ reflects measurement uncertainty or modeling noise. A natural question arises:

710 *What is the possible range of future values $y_1^{(i)}$ under this uncertainty?*

713 A conventional linearized approximation assumes that the system’s response to perturbations is
714 symmetric and smooth around y_0 . Under this assumption, we may apply a first-order Fréchet
715 expansion

$$716 \quad f(y_0 + \Delta y_0, a) \approx f(y_0, a) + \mathcal{D}f(y_0)[\Delta y_0],$$

717 and estimate the extreme values via projection onto the steepest ascent and descent directions.
718 However, for highly nonlinear systems, this linear treatment is inadequate—particularly when the
719 system exhibits directional sensitivity, i.e.,

$$720 \quad f(y_0 + \eta, a) - f(y_0, a) \neq f(y_0 - \eta, a) - f(y_0, a).$$

722 To capture such asymmetry and nonlinear amplification, we adopt a variational formulation. Let the
723 system’s response to a perturbation path η be described by a nonlinear variational functional:

$$724 \quad \mathcal{I}[\eta] = \int_0^1 \mathcal{L}(y_0 + s\eta, \dot{\eta}, s) ds,$$

727 where \mathcal{L} is a nonlinear Lagrangian that encodes the local effect of the perturbation path η and its
728 derivative $\dot{\eta}$ over time or space. The structural asymmetry of the system implies: $\mathcal{I}[\eta] \neq \mathcal{I}[-\eta]$,
729 highlighting that perturbations in opposite directions produce asymmetric responses.

730 We define the extremal range of the future value $y_1^{(i)}$ as the solution to the following constrained
731 nonlinear variational optimization:

$$733 \quad y_1^{(i),\max} = f^{(i)}(y_0) + \sup_{\|\eta\| \leq \Delta y} \left\{ \int_{\mathcal{M}} F^{(i)}[y_0, \eta(x), \nabla \eta(x)] d\mu(x) \right\}, \\ 734 \quad y_1^{(i),\min} = f^{(i)}(y_0) + \inf_{\|\eta\| \leq \Delta y} \left\{ \int_{\mathcal{M}} F^{(i)}[y_0, \eta(x), \nabla \eta(x)] d\mu(x) \right\}, \quad (B.1)$$

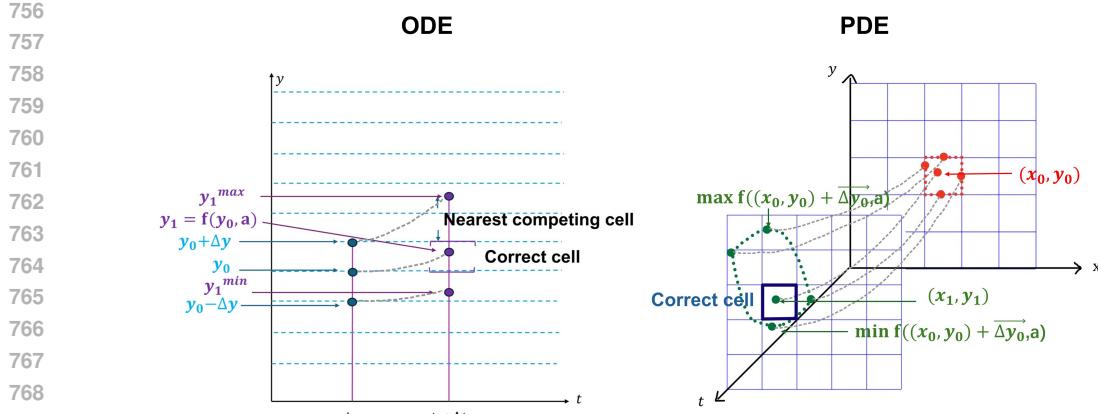
738 where $F^{(i)}$ denotes the nonlinear directional sensitivity of the system output with respect to the i -th
739 component, and $\mu(x)$ is a measure over the spatial domain \mathcal{M} .

740 **Remark B.1.** While this formulation is presented for a single component $y^{(i)}$, it naturally generalizes
741 to the full system state vector. This variational perspective offers a principled and scalable method
742 for quantifying directional uncertainty in nonlinear, asymmetric dynamical systems, with broad
743 applicability to PDE-based forecasting, control, and decision-making.

744 As illustrated in Figure B.1, in the ODE case, the three blue points represent the lower bound, true
745 value, and upper bound of the observed state respectively, resulting from observation noise. Since
746 substituting different observed states into the control equation yields different next-step predicted
747 states, the three purple points correspond to the lower bound, true value, and upper bound of the
748 next-step state.

750 We define a prediction to be correct if it falls into the same grid cell as the true next-step state. This
751 correct cell is represented by the purple box. Predictions that fall into other grid cells are considered
752 incorrect.

753 The PDE example in Figure B.1 illustrates how observation noise in the initial state propagates to
754 the next time step. Here, we show the evolution over a single time step. The red dashed rectangle
755 represents the range of the observed initial state induced by observation noise. The use of a rectangular
region stems from modeling the perturbation δy using the L^∞ norm, which captures worst-case



810 Under the constraint $\|\delta y\|_{L^\infty} \leq \varepsilon$, we then compute bounds:
 811

$$\begin{aligned} 812 \quad \mathcal{F}_{\max} &= \mathcal{F}[y_0] + \sup_{\|\delta y\|_{\infty} \leq \varepsilon} \left\{ \int K(x) \delta y(x) dx + \frac{1}{2} \iint H(x, x') \delta y(x) \delta y(x') dx dx' \right\}, \\ 813 \quad \mathcal{F}_{\min} &= \mathcal{F}[y_0] + \inf_{\|\delta y\|_{\infty} \leq \varepsilon} \left\{ \int K(x) \delta y(x) dx + \frac{1}{2} \iint H(x, x') \delta y(x) \delta y(x') dx dx' \right\}. \end{aligned} \quad (B.2)$$

817 This example demonstrates how second-order variational expansion serves as a concrete realization
 818 of the general nonlinear variational framework. It explicitly captures the directional asymmetry and
 819 curvature of the functional response to perturbations, which are essential features of PDE-governed
 820 systems with uncertain initial conditions.
 821

822 B.2 SAMPLING SCHEME FOR IDENTIFYING THE TRUE DISCRETE CELL

824 In this section, we analyze how a sampling strategy can be designed to identify the correct discrete
 825 cell that the system state should evolve into, and how the scale of such a cell should be chosen.
 826

827 In real-world systems, perturbations Δy typically originate from measurement noise or numerical
 828 rounding errors, and are therefore inevitable. Nevertheless, it is reasonable to assume that the true
 829 state lies within a neighborhood of size $\Delta y/2$ around the observed value. Based on this fact, we
 830 propose a local sampling strategy: perturbations are generated around the observed value within this
 831 neighborhood, and their evolution over one time step is computed via numerical PDE solvers. The
 832 resulting predicted outcomes are then statistically analyzed to infer the most probable cell in which
 833 the true state lies.

834 In the ideal noiseless case, such a sampling scheme causes predicted outcomes to concentrate within
 835 the correct cell. As the number of samples increases, the probability of correctly identifying the target
 836 cell converges to one.

837 This idea naturally supports the integration of physical knowledge into reinforcement learning. By
 838 repeatedly sampling perturbations around the initial state and propagating them through the physical
 839 model, we can estimate the state transition probability matrix with high confidence. This effectively
 840 transforms the reinforcement learning problem from a model-free setting to a dynamic programming
 841 problem with known transitions. In contrast, without such sampling, the system dynamics remain
 842 unknown, and exploration-based algorithms must be employed to estimate both the model and the
 843 optimal policy.

844 To ensure probabilistic determinability of the future state, we impose the following criterion: the
 845 number of predicted samples that fall into the correct cell must exceed that of any incorrect cell.
 846 Specifically, if the predicted uncertainty spans the correct cell and its two immediate neighbors,
 847 and if the correct cell is slightly larger in volume than its neighbors, then the correct cell can be
 848 statistically identified with high confidence. However, if one of the neighboring cells (e.g., the upper
 849 cell) is significantly larger than the others, then it becomes necessary to ensure that the number of
 850 predicted samples falling into the correct cell still exceeds those falling into the largest competing
 851 cell. This places stricter constraints on the acceptable error magnitude Δy , or alternatively, requires
 852 more samples to suppress statistical variance.

853 To improve sampling efficiency, we further propose a variationally-derived optimal importance
 854 sampling strategy. When the goal is to estimate the expected value of a functional output $\mathcal{F}(y_0 + \delta y)$,
 855 importance sampling theory shows that the optimal sampling distribution $p^*(\delta y)$ should satisfy
 856 $p^*(\delta y) \propto |\mathcal{F}(y_0 + \delta y)| \cdot \rho(\delta y)$, where $\rho(\delta y)$ is the base distribution of uncertainty. In the small-
 857 perturbation regime, we expand $\mathcal{F}(y_0 + \delta y)$ using a second-order Taylor approximation:

$$858 \quad \mathcal{F}(y_0 + \delta y) \approx \mathcal{F}(y_0) + \langle \nabla \mathcal{F}, \delta y \rangle + \frac{1}{2} \delta y^\top H \delta y,$$

860 leading to the approximate optimal sampling density $p^*(\delta y) \propto |\langle \nabla \mathcal{F}, \delta y \rangle + \frac{1}{2} \delta y^\top H \delta y|$. Accord-
 861 ingly, we define a sampling weight function: $w(\delta y) := |\langle \nabla \mathcal{F}, \delta y \rangle + \frac{1}{2} \delta y^\top H \delta y|$, which can be used
 862 to guide sampling toward directions with the most significant influence on the predicted outcome.
 863 This minimizes estimation variance and achieves optimal sampling efficiency. The resulting scheme

864 can be interpreted as a physics-informed importance sampling strategy, grounded in variational
 865 sensitivity analysis and importance sampling theory.
 866

867 We now quantify how many samples are needed to reliably identify the correct discrete cell when using
 868 the optimal importance sampling strategy described previously. Specifically, we define the *dominant*
 869 *cell* as the cell that receives the largest number of predicted samples under a given sampling scheme.
 870 Let C^* denote the true (correct) cell and C_j denote any competing cell. Let $\pi^* := \mathbb{P}[\mathcal{F}(y_0 + \delta y) \in C^*]$
 871 and $\pi_j := \mathbb{P}[\mathcal{F}(y_0 + \delta y) \in C_j]$, where $\delta y \sim p^*(\delta y) \propto |\mathcal{F}(y_0 + \delta y)| \cdot \rho(\delta y)$ is the optimal importance
 872 sampling distribution.

873 **Theorem B.1** (Sample Complexity for Correct Cell Dominance). *Let K be the total number of
 874 cells, and let $\delta \in (0, 1)$ be the desired confidence level. Assume that $\Delta := \pi^* - \pi_{\max} > 0$, where
 875 $\pi_{\max} = \max_{j \neq C^*} \pi_j$. Then, in order for the correct cell C^* to receive the largest number of samples
 876 with probability at least $1 - \delta$, it suffices to sample:*

$$877 N \geq \frac{8}{\Delta^2} \cdot \log \left(\frac{2(K-1)}{\delta} \right)$$

879 samples from the optimal importance sampling distribution $p^*(\delta y)$.
 880

881 *Proof.* Let $n_{C^*} \sim \text{Binomial}(N, \pi^*)$ and $n_j \sim \text{Binomial}(N, \pi_j)$ for each $j \neq C^*$. To ensure that
 882 $n_{C^*} > n_j$, we use a union bound and Chernoff-type large deviation inequality. The error probability
 883 that any incorrect cell receives more samples than C^* is bounded as:
 884

$$885 \mathbb{P}(\exists j \neq C^* : n_j \geq n_{C^*}) \leq 2(K-1) \cdot \exp \left(-\frac{N\Delta^2}{8} \right).$$

887 Setting the right-hand side less than or equal to δ yields the desired bound on N . \square
 888

889 This result shows that the sample complexity scales inversely with the square of the gap $\Delta =$
 890 $\pi^* - \pi_{\max}$ between the correct cell and the nearest competitor, and logarithmically with the number
 891 of cells and the inverse failure probability. This quantifies how distinguishable the correct cell is
 892 under the importance sampling distribution. The smaller the margin Δ , the more samples are needed
 893 to overcome statistical uncertainty.
 894

895 However, to preserve generality and narrative consistency, we adopt a uniform sampling strategy in
 896 the main analysis to estimate sample complexity, which differs from the optimal importance sampling
 897 scheme only by a constant-factor overhead.
 898

899 B.3 PROBLEM FORMULATION

900 We consider two distinct scenarios concerning the availability of transition dynamics. These two
 901 cases correspond to two different learning regimes, each with its own complexity characteristics.
 902

903 In the following analysis, we quantify the sample complexity under both settings. Based on
 904 Lemma B.1 and Theorem B.2, we derive two different orders of sample complexity. The bound in
 905 Lemma B.1 captures the general case without prior knowledge of the transition dynamics. In contrast,
 906 Theorem B.2 establishes the bound under our setting, where the transition matrix is known through
 907 repeated sampling from the PDE-based physical space.
 908

909 Therefore, if no sampling is performed in the physical space, the problem remains a standard MDP,
 910 and exploration is required to learn the model. However, if sufficient sampling is carried out in
 911 the PDE domain, the problem effectively becomes a dynamic programming task, and the sample
 912 complexity is driven by the cost of interaction with the physical model, resulting in a tighter bound.
 913 In the main body of the paper, we adopt the sample complexity bound given by Theorem B.2. In
 914 the following, we present several standard assumptions commonly used in the conventional MDP
 915 framework.

916 **Markov Decision Problems** We consider episodic reinforcement learning in finite-horizon MDPs
 917 defined by $\langle S, A, P, R, H \rangle$, where S and A are the finite sets of states and actions, P is the state
 918 transition distribution, the function R is a real-valued function which is deterministic and belongs to
 919 the interval $[0, 1]$, and the horizon H is the length of the episode. We denote by $P(\cdot | x, a)$ and $R(x, a)$

918 the probability distribution over the next state and the immediate reward of taking action a at state x ,
 919 respectively.

920 The agent interacts with the environment in a sequence of episodes. Specifically, the agent follows
 921 a policy $\pi : S \times [H] \rightarrow A$, which maps each state and time step to an action. The value function
 922 $V_h^\pi : S \rightarrow \mathbb{R}$ represents the expected cumulative reward from step h to H under policy π , starting
 923 from state x at time h . Under the previously defined MDP setting, there always exists an optimal
 924 policy π^* attaining the optimal value function $V_h^*(x) \triangleq \sup_\pi V_h^\pi(x)$ for all $x \in S$ and $h \in [H]$.
 925

926 Each policy π induces a transition kernel $P_h^\pi(y|x) \triangleq P(y|x, \pi(x, h))$ and reward function
 927 $r_h^\pi(x) \triangleq R(x, \pi(x, h))$. For any function $V : S \rightarrow \mathbb{R}$, we define the transition operators:
 928 $(P_h^\pi V)(x) \triangleq \sum_{y \in S} P_h^\pi(y|x)V(y)$. The Bellman operator for policy π is defined as $(\mathcal{T}_h^\pi V)(x) \triangleq$
 929 $r_h^\pi(x) + (P_h^\pi V)(x)$.
 930

931 We assume that the reward function r is known to the agent but the transition kernel p is unknown
 932 (Under the no-sampling setting). The question we study is how many episodes does a learning agent
 933 follow a policy π that is not ϵ -optimal, i.e., $V_M^* - \epsilon > V_M^\pi$, with probability at least $1 - \delta$ for any
 934 chosen accuracy ϵ and failure probability δ .
 935

B.4 FULL ANALYSIS OF SAMPLE COMPLEXITY BOUND

937 In this section, we systematically analyze the PAC sample complexity of reinforcement learning
 938 under two different assumptions regarding the transition probability matrix. We begin by examining
 939 the classical setting where the transition dynamics are entirely unknown. Lemma B.1 provides a
 940 baseline PAC bound in this setting, and reveal how the total sample complexity arises from key error
 941 components, including estimation error, exploration bonuses, and confidence adjustments.

942 Building on this foundation, we then analyze how our proposed setting—where additional structure
 943 or prior knowledge is introduced—modifies the learning process. In particular, Lemmas B.2 and
 944 B.3 show how specific components in the original analysis are tightened or avoided, leading to an
 945 improved sample complexity bound.

946 **Theorem B.2** (PAC Bound under our setting). *Assume the transition dynamics of the MDP are
 947 known up to statistical confidence level δ , reducing the reinforcement learning problem to dynamic
 948 programming. Let the correct classification cell have predicted frequency $p = \Delta y^{(j)}$, and let the
 949 most probable competing cell have frequency $q = p_{\max}^{(j)}$. Then the number of forward predictions
 950 required to identify the correct classification cell with probability at least $1 - \delta$ satisfies*

$$951 \quad 952 \quad 953 \quad 954 \quad n = \mathcal{O} \left(HSA \cdot \frac{\log(1/\delta)}{\min_j \left(\Delta y^{(j)} - p_{\max}^{(j)} \right)^2} \right).$$

955 Furthermore, we derive a refined bound on the sample complexity $n = \mathcal{O} \left(\frac{SAH^3}{\epsilon^2} \cdot \log(1/\delta) \right)$.
 956

957 We also present in Lemma B.1 the PAC bound of the classical UCB-VI algorithm under the setting
 958 where the state transition probability matrix is entirely unknown. The detailed algorithmic procedure
 959 of UCB-VI can be found in (Azar et al., 2017).
 960

961 **Lemma B.1.** *For any $0 < \epsilon, \delta \leq 1$, the following holds. With probability at least $1 - \delta$, the algorithm
 962 UCB – VI produces a sequence of policies π^k , that yield at most $\mathcal{O} \left(\frac{SAH^3}{\epsilon^2} \log \frac{1}{\delta} \right)$ episodes where
 963 $V_{1:H}^*(s_0) - V_{1:H}^{\pi^k}(s_0) > \epsilon$. This results in a total runtime for sampling of $\mathcal{O} \left(\frac{SAH^4}{\epsilon^2} \log \frac{1}{\delta} \right)$.*
 964

965 *Proof.* Lemma B.1 characterizes the PAC bound of the UCB-VI algorithm, a result that has been
 966 rigorously established and thoroughly analyzed in prior work (see Azar et al. (2017)). To facilitate
 967 our later analysis, we provide a concise overview of the proof from Azar et al. (2017), with particular
 968 emphasis on the error decomposition framework and the core techniques used to bound individual
 969 error components.
 970

971 Let $\mathcal{V} = \{V_{k,h} \geq V_h^*, \forall k, h\}$ denote the event where all computed value estimates are optimistic
 972 upper bounds of the optimal value function. Azar et al. (2017) show that \mathcal{V} holds with probability

at least $1 - \delta$ under standard induction arguments with appropriate bonus functions. Define the approximation errors $w_{k,h} = W_{k,h}(x_{k,h})$ and $\tilde{w}_{k,h} = \tilde{W}_{k,h}(x_{k,h})$, where $W_{k,h} = V_h^* - V_h^{\pi_k}$ and $\tilde{W}_{k,h} = V_{k,h} - V_h^{\pi_k}$. Then, under the event \mathcal{V} , the following error decomposition holds:

$$w_{k,h} \leq \tilde{w}_{k,h} = \mathcal{O} \left(\sum_{i=h}^{H-1} \left(M_{k,i} + 2\sqrt{L\bar{M}_{k,i}} + \text{bonus}_{k,i} + E_{k,i} \right) \right),$$

where $M_{k,h} = P^{\pi_k}W_{k,h+1}(x_{k,h}) - W_{k,h+1}(x_{k,h+1})$ is a martingale difference term. $\bar{M}_{k,i}$ is the quantity resulting from applying Bernstein's inequality to control the term $[(\hat{P}_h^{\pi_k} - P_h^{\pi_k})(V_{h+1} - V_{h+1}^*)]_{(x_h)}$ and is also a martingale difference. $\text{bonus}_{k,i}$ represents the bonus function for UCB-VI at step i under episode k . $E_{k,h} \stackrel{\text{def}}{=} (\hat{P}_k^{\pi_k} - P_k^{\pi_k})V_{h+1}^*(x_{k,h})$ is the estimation error of the optimal value function at the next state. In addition, L here is a logarithmic factor depending on H, S, A , and the confidence level δ , which does not affect the polynomial complexity of the bound.

Based on the above statement, we observe that the UCB-VI algorithm decomposes the total error into four components: $M_{k,i}$, $\bar{M}_{k,i}$, $\text{bonus}_{k,i}$, and $E_{k,i}$. In Table B.2, we summarize the order of magnitude and the bounding techniques for each of these four error terms. The detailed proof process can be found in (Azar et al., 2017). Therefore, the cumulative value error over all episodes and timesteps satisfies $\sum_{k,h} w_{k,h} = \mathcal{O}(HL\sqrt{SAT})$, where $T = KH$ is the total number of steps collected across K episodes, each of length H . This implies:

$$\sum_{k,h} w_{k,h} = \mathcal{O}(HL\sqrt{SAKH}) = \mathcal{O}(H^{3/2}L\sqrt{SAK}).$$

Suppose in the worst-case scenario, all of the first K episodes are not ε -optimal before an ε -optimal policy is learned, then we require:

$$K_{\text{bad}} \cdot \varepsilon \leq \sum_{k,h} w_{k,h} = \mathcal{O}(H^{3/2}L\sqrt{SAK_{\text{bad}}}) \Rightarrow K_{\text{bad}} \leq \mathcal{O}\left(\frac{H^3L^2SA}{\varepsilon^2}\right) = \mathcal{O}\left(\frac{SAH^3}{\varepsilon^2} \log \frac{1}{\delta}\right).$$

To derive the total sample complexity, we note that sampling one episode and updating the respective variables has $\mathcal{O}(H)$ runtime. Thus, the total runtime across all K_{bad} episodes is $\mathcal{O}(H \cdot K_{\text{bad}}) = \mathcal{O}\left(\frac{SAH^4}{\varepsilon^2} \log \frac{1}{\delta}\right)$. This establishes the total sample complexity required to ensure that the learned policy is ε -optimal with high probability.

Table B.2: Error decomposition in UCB-VI and bounding techniques

Error Term	Magnitude	Bounding Technique
$\sum_{k,h} M_{k,h}$	$\mathcal{O}(H\sqrt{TL})$	Azuma's inequality (martingale bound)
$\sum_{k,h} \bar{M}_{k,h}$	$\mathcal{O}(\sqrt{TL})$	Azuma's inequality (martingale bound)
$\sum_{k,h} E_{k,h}$	$\mathcal{O}(H\sqrt{LSAT})$	Chernoff-Hoeffding inequality + pigeonhole principle
$\sum_{k,h} \text{bonus}_{k,h}$	$\mathcal{O}(HL\sqrt{SAT})$	Optimistic bonus design + pigeonhole principle

□

Remark B.2. Based on Table B.2, we have essentially clarified the analytical procedure through which the UCB-VI algorithm derives its sample complexity bound. Based on this, we highlight the key difference between our setting and UCB-VI, and the core improvement introduced in our approach.

UCB-VI characterizes sample complexity by bounding the discrepancy between the empirical transition probability matrix obtained from sampling and the true transition dynamics.

1026
 1027 *In contrast, our setting introduces a notion of the probabilistic determinability of future state, which*
 1028 *allows us to exploit multiple samples from the PDE-based physical space to ensure that the transition*
 1029 *probability matrix in the RL formulation is fully known.*

1030 *As a result, the increase in sample size in our framework pertains to the sampling process in the*
 1031 *PDE-based physical space. The complexity of this sampling process can then be transferred to the*
 1032 *RL domain, ultimately yielding a refined bound on sample complexity.*

1033 In Lemma B.2, we provide the sampling requirement needed to resolve the correct cell with probability
 1034 at least $1 - \delta$.

1035 **Lemma B.2** (Sample Complexity for δ -Confidence Classification). *Let the cell of correct classifica-*
 1036 *tion have predicted frequency $p = \Delta y^{(j)}$, and let the most probable incorrect (competing) cell have*
 1037 *frequency $q = p_{\max}^{(j)}$. To resolve the correct cell with confidence at least $1 - \delta$, the number of forward*
 1038 *predictions required satisfies*

$$1040 \quad 1041 \quad 1042 \quad 1043 \quad n = \mathcal{O} \left(\frac{\log(1/\delta)}{\min_j (\Delta y^{(j)} - p_{\max}^{(j)})^2} \right).$$

1044
 1045 *Proof.* Without loss of generality, we consider a randomly selected dimension j as an example. Let
 1046 $p = \Delta y^{(j)}$ be the probability mass of the correct cell in dimension j , and let $q = p_{\max}^{(j)}$ be the
 1047 maximum probability mass of any competing cell.

1048 Suppose we sample n i.i.d. trajectories, and let \hat{p} denotes the empirical frequencies of the correct
 1049 cell. We wish to ensure that with high probability (at least $1 - \delta$), $|\hat{p} - p| < \frac{|p - q|}{2}$, so that the correct
 1050 classification cell is selected.

1051 Let $\epsilon = \frac{|p - q|}{2}$. Applying Hoeffding's inequality to \hat{p} gives:

$$1054 \quad 1055 \quad \Pr (|\hat{p} - p| \geq \epsilon) \leq 2 \exp(-2n\epsilon^2) \Rightarrow \Pr \left(|\hat{p} - p| \geq \frac{|p - q|}{2} \right) \leq 2 \exp(-2n(\frac{|p - q|}{2})^2).$$

1056
 1057 Then enforce $2 \exp \left(-2n \left(\frac{p - q}{2} \right)^2 \right) \leq \delta \Rightarrow n \geq \frac{2 \log(2/\delta)}{(p - q)^2}$.

1058 This gives the number of samples n needed to distinguish between the correct and competing cells
 1059 with probability at least $1 - \delta$. To guarantee correct classification across all dimensions j , we must
 1060 take the worst-case (smallest gap) over all j , yielding: $n = \mathcal{O} \left(\frac{\log(1/\delta)}{\min_j (\Delta y^{(j)} - p_{\max}^{(j)})^2} \right)$ as stated. \square

1061 Furthermore, Lemma B.3 provides the order of the total sample complexity under our setting.

1062 **Lemma B.3** (Refined Sample Complexity under Known Transitions). *Assume the transition dynamics*
 1063 *of the MDP are known up to statistical confidence level δ , reducing the reinforcement learning problem*
 1064 *to dynamic programming. Let the correct classification cell have predicted frequency $p = \Delta y^{(j)}$,*
 1065 *and let the most probable competing cell have frequency $q = p_{\max}^{(j)}$. Then the number of forward*
 1066 *predictions required to identify the correct classification cell with probability at least $1 - \delta$ satisfies*

$$1067 \quad 1068 \quad 1069 \quad 1070 \quad n = \mathcal{O} \left(H_r S_r A_r \cdot \frac{\log(1/\delta)}{\min_j (\Delta y^{(j)} - p_{\max}^{(j)})^2} \right).$$

1071 Furthermore, by applying the $\rho - K$ relationship discussed in the main text, which states that
 1072 $\Delta p^{(j)} \sim \rho \Delta y^{(j)} \sim \frac{1}{H_r} \Delta y^{(j)}$, we derive a refined bound on the sample complexity:

$$1073 \quad 1074 \quad 1075 \quad n = \mathcal{O} (H_r^3 S_r A_r \cdot \log(1/\delta)).$$

1080 *Proof.* Without loss of generality, consider a randomly selected dimension j . Let $p = \Delta y^{(j)}$ be the
 1081 predicted frequency of the correct classification cell, and let $q = p_{\max}^{(j)}$ denote the predicted frequency
 1082 of the most likely competing cell in dimension j .
 1083

1084 Since the transition matrix is assumed to be known up to confidence δ , the reinforcement learning
 1085 task reduces to a dynamic programming problem. Assume we sample n independent forward
 1086 trajectories and let \hat{p} be the empirical frequency of the correct cell. To confidently resolve the correct
 1087 classification, we want the empirical estimate to satisfy $|\hat{p} - p| < \frac{|p - q|}{2}$ with probability at least $1 - \delta$.
 1088 Let $\epsilon = \frac{|p - q|}{2}$. According to Lemma B.2, we have $n \geq \frac{2 \log(2/\delta)}{(p - q)^2}$.
 1089

1090 At each step in the reinforcement learning process, every state-action pair requires the number of
 1091 samples derived above to guarantee correct classification with high confidence. Applying a union
 1092 bound over all such points yields a total sample complexity of:
 1093

$$n = \mathcal{O} \left(H_r S_r A_r \cdot \frac{\log(1/\delta)}{\min_j (p - q)^2} \right).$$

1096 Furthermore, we have: $p - q \sim \mathcal{O} \left(\frac{1}{H_r} \right)$, thus $\frac{1}{(p - q)^2} = \mathcal{O}(H_r^2)$. Substituting this into the
 1097 earlier bound yields: $n = \mathcal{O} (H_r^3 S_r A_r \cdot \log(1/\delta))$ as claimed. \square
 1098

1100 **Remark B.3.** *Based on Lemma B.1 and Theorem . B.2, we derive two different orders of sample
 1101 complexity. However, we also leave open the possibility for further refinement. In particular, building
 1102 upon the analytical framework presented above, future research may explore combining our treatment
 1103 of transition uncertainty with recent developments on PAC bounds in robust MDPs, potentially
 1104 yielding tighter complexity guarantees.*

1107 C DETAILED ANALYSIS OF THE FOUR REPRESENTATIVE AI4S-RL 1108 ENVIRONMENTS

1110 In this section, building on the theoretical framework proposed in the main text, we provide a
 1111 detailed analysis of how surrogate discretization, spectral response, and reinforcement learning
 1112 precision jointly affect sample complexity across four representative AI4S-RL systems: (1) Tokamak
 1113 plasma control, (2) Turbulent airfoil regulation, (3) Teppanyaki heat sequencing, and (4) Cart-pole
 1114 stabilization. The first three environments are PDE-based systems, while the fourth represents
 1115 a typical ODE-based system. Notably, the first two environments are characterized as systems
 1116 with strong boundary observability. Therefore, by conducting a detailed discussion of these four
 1117 environments, we analyze how different environmental characteristics and underlying equation types
 1118 influence the outcomes and properties of optimal resolution matching and computational cost, thereby
 1119 enhancing the generality of our conclusions.

1121 C.1 THEORETICAL PRELIMINARIES AND PROOFS

1123 Before proceeding to the detailed analysis of each environment, this subsection provides several
 1124 general settings and supplementary explanations for key lemmas and conclusions referenced in the
 1125 main text. In addition, Appendix D provides a collection of classical results from functional analysis
 1126 and PDE theory that underpin the error analysis.

1128 C.1.1 PROOF OF THEOREM 1

1130 **Theorem C.3** (Sample Complexity for δ -Confidence Classification). *Consider repeated forward
 1131 predictions from perturbed initial states, where the predicted frequency is the empirical probability of
 1132 the next state falling into a given grid cell. Let p denote the predicted frequency with which the true
 1133 next state falls into the correct grid cell, and let $q = p_{\max}^{(j)}$ denote the maximum predicted frequency
 among all competing cells. To resolve the correct cell with confidence at least $1 - \delta$, the number of*

1134 forward predictions required is bounded by
 1135

$$1136 \quad 1137 \quad 1138 \quad 1139 \quad n = \mathcal{O} \left(\frac{\log(1/\delta)}{\min_j (\Delta y^{(j)} - p_{\max}^{(j)})^2} \right) \quad (C.1)$$

1140 *Proof.* We proceed under the fundamental assumption of **state distinguishability**, which ensures that
 1141 the propagation of initial uncertainty is sufficiently bounded. This allows for the correct classification
 1142 cell to be identified, as repeated sampling from the input perturbation ball will generate predicted
 1143 states that form a distinct cluster primarily within that cell.

1144 The problem can thus be framed as a statistical task: we perform n independent trials (forward
 1145 predictions) to identify the cell with the highest success probability, p , from its closest competitor,
 1146 which has probability q . Let X_i be the indicator random variable for the i -th prediction falling
 1147 into the correct cell. The empirical frequency of the correct cell, $\hat{p} = \frac{1}{n} \sum_{i=1}^n X_i$, is the sample
 1148 mean of these variables, with an expected value Consider n independent forward predictions, each
 1149 obtained by sampling a perturbation Δy_0 uniformly from the ball $\{\|\Delta y_0\| \leq \Delta y\}$ and computing
 1150 $\tilde{y}_1 = f(y_0 + \Delta y_0, a)$. Let X_i be the indicator random variable for the i -th prediction falling into the
 1151 correct cell: $X_i = \begin{cases} 1 & \text{if } \tilde{y}_1^{(i)} \text{ falls in the correct cell} \\ 0 & \text{otherwise} \end{cases}$. The empirical frequency of the correct cell is:

1152 $\hat{p} = \frac{1}{n} \sum_{i=1}^n X_i$. By construction, $\mathbb{E}[X_i] = p$ and $X_i \in [0, 1]$.

1153 Since X_1, X_2, \dots, X_n are independent and bounded random variables with $X_i \in [0, 1]$, Hoeffding's
 1154 inequality states: $\Pr(|\hat{p} - p| \geq \epsilon) \leq 2 \exp(-2n\epsilon^2)$ for any $\epsilon > 0$. To distinguish the correct cell
 1155 from competing cells, we need \hat{p} to be closer to p than to q . This is achieved if: $|\hat{p} - p| < \frac{|p-q|}{2}$.
 1156 Setting $\epsilon = \frac{|p-q|}{2} = \frac{p-q}{2}$ (since $p > q$ by assumption), we get:

$$1157 \quad 1158 \quad 1159 \quad \Pr \left(|\hat{p} - p| \geq \frac{p - q}{2} \right) \leq 2 \exp \left(-2n \left(\frac{p - q}{2} \right)^2 \right)$$

1160 For the classification to succeed with confidence at least $1 - \delta$, we require:

$$1161 \quad 1162 \quad 1163 \quad \Pr \left(|\hat{p} - p| \geq \frac{p - q}{2} \right) \leq 2 \exp \left(-2n \left(\frac{p - q}{2} \right)^2 \right) \leq \delta$$

1164 Taking logarithms and solving for n : $n \geq \frac{2 \log(2/\delta)}{(p-q)^2}$. When multiple dimensions are involved, we
 1165 need to ensure correct classification across all dimensions j . Taking the worst-case scenario (smallest
 1166 gap):

$$1167 \quad 1168 \quad 1169 \quad n \geq \max_j \frac{2 \log(2/\delta)}{(\Delta y^{(j)} - p_{\max}^{(j)})^2} = \frac{2 \log(2/\delta)}{\min_j (\Delta y^{(j)} - p_{\max}^{(j)})^2}$$

1170 In terms of asymptotic complexity (Big-O notation), the constant factors can be absorbed, yielding
 1171 the final result: $n = \mathcal{O} \left(\frac{\log(1/\delta)}{\min_j (\Delta y^{(j)} - p_{\max}^{(j)})^2} \right)$. This completes the proof of the sample complexity
 1172 bound. \square

1173

1174 C.1.2 MODAL GROWTH RATE AND STATE SEPARABILITY

1175

1176 In AI4S-RL, agents interact with high-dimensional physical systems, where observations are often
 1177 contaminated with sensor noise or state perturbations. A key concern is whether such uncertainty
 1178 grow through the system's nonlinear dynamics and cause classification ambiguity in the predicted
 1179 future state.

1180

1181 To ensure robustness of the learning process, we must guarantee that small observation errors do
 1182 not cause the predicted future state $f(y_0 + \Delta y_0)$ to cross classification boundaries. This leads to the
 1183 formulation of Lemma S.3, which characterizes sufficient conditions for achieving classification-cell
 1184 separability.

1185

1186

1187

1188
 1189 **Lemma S.3** (Fourier-Mode Separability via Fréchet Spectral Amplification). *Let $f_t(y_0)$ be the*
 1190 *solution operator of a nonlinear PDE with Fréchet derivative $\mathcal{D}f_t(y_0)$ acting on a Hilbert space*
 1191 *\mathcal{H} . Suppose the observation perturbation $\eta \in \mathcal{H}$, with $\|\eta\| \leq \Delta y$, admits a Fourier expansion*
 1192 *$\eta = \sum_k \hat{\eta}_k \phi_k$, where $\{\phi_k\}$ is an orthonormal basis for \mathcal{H} .*

1193 *Assume the Fréchet derivative propagates each mode as $\mathcal{D}f_t(y_0)[\phi_k] = \gamma_k(t) \psi_k(t)$ for some gain*
 1194 *factor $\gamma_k(t) > 0$. Then classification-cell separability is guaranteed if: $\sup_k \gamma_k(t) < 1$, which*
 1195 *implies $\|\mathcal{D}f_t(y_0)[\eta]\| < \Delta y$. Consequently, the maximal admissible time step Δt satisfies:*

$$1196 \quad \Delta t < C_{\text{system}} := \inf_k \gamma_k^{-1}(1).$$

1198 *Proof.* Let $\eta = \sum_k \hat{\eta}_k \phi_k \in \mathcal{H}$ be the perturbation with norm $\|\eta\| \leq \Delta y$, where $\{\phi_k\}$ is an
 1199 orthonormal basis. Since $\mathcal{D}f_t(y_0)$ is linear, then we have:

$$1201 \quad \mathcal{D}f_t(y_0)[\eta] = \sum_k \hat{\eta}_k \cdot \mathcal{D}f_t(y_0)[\phi_k] = \sum_k \hat{\eta}_k \cdot \gamma_k(t) \cdot \psi_k(t).$$

1203 Now take the squared norm:

$$1205 \quad \|\mathcal{D}f_t(y_0)[\eta]\|^2 = \left\| \sum_k \hat{\eta}_k \cdot \gamma_k(t) \right\|^2 \leq \left(\sup_k \gamma_k(t)^2 \right) \cdot \sum_k |\hat{\eta}_k|^2 = \sup_k \gamma_k(t)^2 \cdot \|\eta\|^2.$$

1208 Therefore,

$$1210 \quad \|\mathcal{D}f_t(y_0)[\eta]\| \leq \sup_k \gamma_k(t) \cdot \|\eta\| \leq \sup_k \gamma_k(t) \cdot \Delta y.$$

1211 To ensure $\|\mathcal{D}f_t(y_0)[\eta]\| < \Delta y$, it suffices that $\sup_k \gamma_k(t) < 1$. \square

1214 The spectral separability condition established in Lemma S.3 applies broadly to a wide class of
 1215 PDE-governed systems. In general, when the system dynamics are governed by a nonlinear evolution
 1216 equation, local behavior around state y_0 can be analyzed through Fréchet linearization, yielding a
 1217 time-dependent linear operator $\mathcal{D}f_t(y_0)$. By decomposing initial observation perturbations $\eta \in \mathcal{H}$
 1218 into orthonormal modal components and tracking the amplification of each mode under the action of
 1219 $\mathcal{D}f_t(y_0)$, we obtain a spectral characterization of uncertainty propagation over time.

1220 C.1.3 RELATING MODAL GAIN $\gamma_k(t)$ TO THE LEADING EIGENVALUE λ_1

1222 In the framework discussed in the previous section, each mode experiences a gain factor $\gamma_k(t)$,
 1223 which reflects the sensitivity of the system to perturbations in that direction. When the maximal gain
 1224 $\sup_k \gamma_k(t)$ remains below unity, the propagated perturbation stays within the classification boundary,
 1225 ensuring separability and decision consistency. This condition naturally induces a system-dependent
 1226 constraint on temporal resolution: the RL decision step size must be sufficiently high to resolve the
 1227 most unstable direction in the spectral space. In practice, this analysis enables principled derivation
 1228 of time-step bounds for a variety of AI4S-RL environments, thereby enabling robust policy learning
 1229 grounded in physical dynamics.

1230 **Remark C.4.** *In nonlinear dynamical systems governed by PDEs, the rigorous characterization of*
 1231 *forward perturbation growth should ideally be expressed in terms of the Fréchet-mode amplification*
 1232 *factors $\gamma_k(t)$, as introduced in Lemma S.3. These quantities measure the directional sensitivity of the*
 1233 *solution operator f_t along each Fourier mode ϕ_k , through the relation $\mathcal{D}f_t(y_0)[\phi_k] = \gamma_k(t) \psi_k(t)$,*
 1234 *where $\psi_k(t)$ denotes the propagated mode shape. However, to simplify the analysis and maintain*
 1235 *consistency across sections, we adopt the leading eigenvalue λ_1 of the linearized system operator*
 1236 *$\mathcal{L}(y_0)$ as a surrogate for modal growth, particularly in the computation of time resolution bounds*
 1237 *and error propagation rates. The relationship between λ_k and $\gamma_k(t)$ can be made explicit under the*
 1238 *assumption of locally linearized evolution:*

$$1238 \quad \delta y_k(t) = \hat{\eta}_k e^{\lambda_k t} \phi_k \quad \Rightarrow \quad \gamma_k(t) \propto e^{\lambda_k t}.$$

1240 For sufficiently small t , Taylor expansion yields: $\gamma_k(t) \approx 1 + \lambda_k t + \mathcal{O}(t^2)$, so that $\sup_k \gamma_k(t) \propto \lambda_1$.
 1241 Therefore, replacing $\sup_k \gamma_k(t)$ with λ_1 in our main analysis introduces no change in asymptotic
 1242 error scaling.

1242 C.1.4 PROOF OF THEOREM 2
12431244 **Theorem C.4** (ρ -K Analysis for d -Dimensional Systems). *For a d -dimensional PDE system, the*
1245 *forward projection error rate is:*

1246
1247
$$\rho = 1 - \frac{1}{\lambda_1/H_r + 1 + C_1/K^d} \quad (\text{C.2})$$

1248

1249 *where the numerical error scales as $\Delta_{\text{total}} \sim C_1 K^{-d} \Delta y$ from spatial discretization. In the high-*
1250 *resolution limit:*

1251
1252
$$\rho = \mathcal{O}\left(\frac{1}{H_r} + \frac{1}{K^d}\right) \quad (\text{C.3})$$

1253

1254
1255 *Proof.* The forward projection error rate, ρ , is defined as the fractional deviation from a perfect
1256 prediction. Conceptually, this is one minus the ratio of the ideal state-cell size to the total spread of
1257 the predicted state under all sources of error. Let the ideal cell size be represented by the baseline
1258 observation uncertainty, Δy .

1259
1260
$$\rho = 1 - \frac{\text{Ideal Cell Size}}{\text{Total Error Spread}} = 1 - \frac{\Delta y}{\text{Total Error Spread}}$$

1261

1262 The total error spread is the linear superposition of three primary components identified in the text:
12631264
1265 1. **Intrinsic Error Growth** ($\Delta_{\text{intrinsic}}$): The amplification of the initial uncertainty Δy by the
1266 system's dynamics over a single RL time step $\Delta t_r = 1/H_r$. This is governed by the leading
1267 eigenvalue λ_1 , yielding $\Delta_{\text{intrinsic}} = \frac{\lambda_1}{H_r} \Delta y$.
1268
1269 2. **Baseline Observation Uncertainty** (Δ_{obs}): The inherent noise floor of the system, which is Δy .
1270
1271 3. **Numerical Surrogate Error** (Δ_{num}): The total error from discretization, Δ_{total} , which the theorem
1272 states scales as $\Delta_{\text{total}} \sim C_1 K^{-d} \Delta y$. This scaling arises from the resolution matching conditions,
1273 where the dominant contribution to error that depends on the resolution ratio K is the spatial
1274 discretization of the underlying PDE in a d -dimensional space.1275 Summing these components and canceling the Δy term yields the first result of the theorem:
1276

1277
1278
$$\rho = 1 - \frac{1}{\frac{\lambda_1}{H_r} + 1 + C_1 K^{-d}}$$

1279

1280 The high-resolution limit is defined as the case where $H_r \rightarrow \infty$ and $K \rightarrow \infty$. To analyze the
1281 behavior of ρ in this limit, we first rearrange the expression:
1282

1283
1284
$$1 - \rho = \frac{1}{1 + \frac{\lambda_1}{H_r} + C_1 K^{-d}}$$

1285

1286 Let $x = \frac{\lambda_1}{H_r} + C_1 K^{-d}$. In the high-resolution limit, as $H_r \rightarrow \infty$ and $K \rightarrow \infty$, it is clear that
1287 $x \rightarrow 0$. We can therefore use the first-order Taylor series expansion for the function $f(x) = \frac{1}{1+x}$
1288 around $x = 0$. Substituting our expression for x : $1 - \rho \approx 1 - \left(\frac{\lambda_1}{H_r} + C_1 K^{-d}\right)$. Then we obtain the
1289 approximate expression for ρ : $\rho \approx \frac{\lambda_1}{H_r} + C_1 K^{-d}$. Since λ_1 and C_1 are constants for a given system,
1290 the asymptotic behavior of ρ is captured by Big-O notation as:
1291

1292
1293
$$\rho = \mathcal{O}\left(\frac{1}{H_r} + \frac{1}{K^d}\right)$$

1294

1295 This completes the proof of the second part of the theorem. \square

1296 C.1.5 PROOF OF THEOREM 3
12971298 **Theorem C.5** (Optimal Resolution with System-Dependent Scaling). *For a physical system with*
1299 *state space scaling $S_r \sim K^\alpha$ and action space scaling $A_r \sim K^\beta$, under the computational balance*
1300 *condition $H_r^3 S_r A_r \sim H_w S_w A_w$, minimizing the computational cost*

1301
1302
$$\text{Cost}(K) = H_r^3 K^{\alpha+\beta} \cdot \left(\frac{\log(1/\delta)}{\varepsilon^2} \right) \cdot \left(\frac{1}{1 - \frac{1}{H_r} - \frac{1}{K^d}} \right)^2 \quad (\text{C.4})$$

1303
1304

1305 *yields optimal resolution:*

1306
1307
$$K^* = \left(\frac{\alpha + \beta + 2d}{(\alpha + \beta)(1 - H_r^{-1})} \right)^{1/d} \approx \left(\frac{\alpha + \beta + 2d}{\alpha + \beta} \right)^{1/d} \cdot \exp \left(\frac{1}{d\lambda_1} \right) \quad (\text{C.5})$$

1308
1309

1310 *when $H_r \gtrsim \lambda_1 \gg 1$.*
13111312 *Proof.* The total computational cost is given by:
1313

1314
1315
$$\text{Cost}(K) = H_r^3 K^{\alpha+\beta} \cdot \frac{\log(1/\delta)}{\varepsilon^2} \cdot \left(\frac{1}{1 - \frac{1}{H_r} - \frac{1}{K^d}} \right)^2$$

1316
1317

1318 To find the minimum, we differentiate with respect to K . Since the prefactor $H_r^3 \frac{\log(1/\delta)}{\varepsilon^2}$ is inde-
1319 pendent of K , we focus on $g(K) = K^{\alpha+\beta} \cdot \left(1 - \frac{1}{H_r} - \frac{1}{K^d}\right)^2$. To simplify the calculation, we use
1320 logarithmic differentiation $\ln g(K) = (\alpha + \beta) \ln K + 2 \ln \left(1 - \frac{1}{H_r} - \frac{1}{K^d}\right)$.
13211322 Differentiating both sides with respect to K : $\frac{g'(K)}{g(K)} = \frac{\alpha + \beta}{K} - \frac{2d}{K^{d+1} \left(1 - \frac{1}{H_r} - \frac{1}{K^d}\right)}$. At the optimal point
1323 K^* , we have $g'(K^*) = 0$, which gives:
1324

1325
1326
$$\alpha + \beta = \frac{2d + \alpha + \beta}{(K^*)^d \left(1 - \frac{1}{H_r}\right)}$$

1327
1328

1329 Then $(K^*)^d = \frac{\alpha + \beta + 2d}{(\alpha + \beta)(1 - H_r^{-1})}$. Taking the d -th root $K^* = \left(\frac{\alpha + \beta + 2d}{(\alpha + \beta)(1 - H_r^{-1})} \right)^{1/d}$.
13301331 We now analyze the expression for K^* in the limit where $H_r \gtrsim \lambda_1 \gg 1$. This implies that $x = H_r^{-1}$
1332 is a small positive value. We can rewrite the expression for K^* as:
1333

1334
1335
$$K^* = \left(\frac{2d + \alpha + \beta}{\alpha + \beta} \right)^{1/d} \cdot (1 - H_r^{-1})^{-1/d}$$

1336
1337

1338 We focus on the second term, $(1 - H_r^{-1})^{-1/d}$. Using the generalized binomial approximation,
1339 $(1 - x)^a \approx 1 - ax$ for small x , where $x = H_r^{-1}$ and $a = -1/d$: $(1 - H_r^{-1})^{-1/d} \approx 1 + \frac{1}{dH_r}$.
1340 For large λ_1 , the argument of the exponential is small. Using the first-order Taylor approximation
1341 $e^y \approx 1 + y$ for small y : $(1 - H_r^{-1})^{-1/d} \approx \exp \left(\frac{1}{d\lambda_1} \right)$. Substituting this back gives the final
1342 approximate form for the optimal resolution:
1343

1344
1345
$$K^* \approx \left(\frac{2d + \alpha + \beta}{\alpha + \beta} \right)^{1/d} \cdot \exp \left(\frac{1}{d\lambda_1} \right)$$

1346
1347

1348 This completes the proof.
1349

□

1350 C.2 TOKAMAK PLASMA CONTROL
13511352 C.2.1 ENVIRONMENT DESCRIPTION
13531354 In this task, we study a magnetically confined plasma system involving the control of a droplet-like
1355 plasma structure within a tokamak device. The objective is to manipulate the external magnetic field
1356 to maintain stable confinement of the plasma droplet and prevent wall contact (Degrave et al., 2022).1357 The governing dynamics are modeled by a coupled magnetohydrodynamic system with deformable
1358 interface dynamics, comprising the following equations:

1359

1360
$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \quad (\text{Continuous equation})$$

1361

1362
$$\rho \left(\frac{\partial \mathbf{v}}{\partial t} + \mathbf{v} \cdot \nabla \mathbf{v} \right) = -\nabla p + \mathbf{J} \times \mathbf{B} + \mu \nabla^2 \mathbf{v} + \mathbf{F}_{\text{surface}} \quad (\text{Momentum equation})$$

1363

1364
$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times (\mathbf{v} \times \mathbf{B}) + \eta \nabla^2 \mathbf{B} \quad (\text{Magnetic field control equation})$$

1365

1366
$$\frac{\partial \phi}{\partial t} + \mathbf{v} \cdot \nabla \phi = 0 \quad (\text{Deformation of the droplet's interface})$$

1367

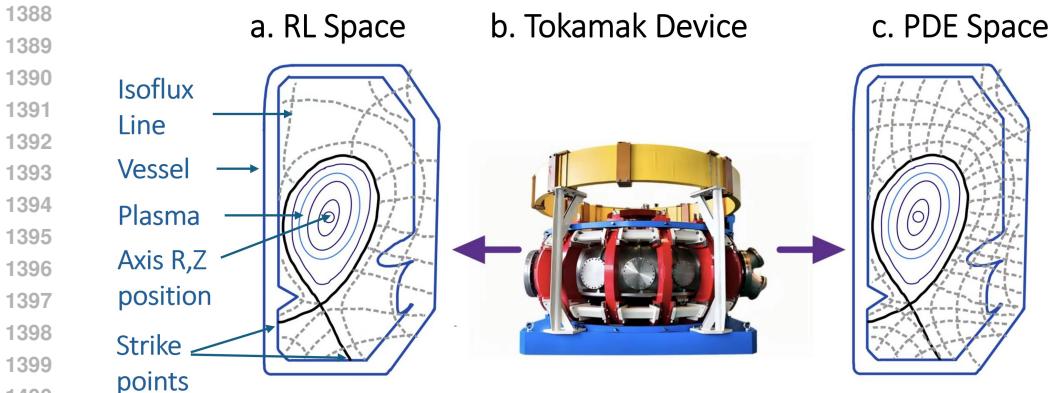
1368
$$\sigma \kappa = \gamma_{\text{surface}} \quad (\text{Boundary conditions})$$

1369

1370

1371 Here, ρ is the plasma density, \mathbf{v} is the velocity field, p is the pressure, \mathbf{B} is the magnetic field, and
1372 $\mathbf{J} = \nabla \times \mathbf{B}$ is the current density. The term μ is the dynamic viscosity, η is the magnetic resistivity,
1373 and $\mathbf{F}_{\text{surface}}$ represents interface forces. The function ϕ is a level-set representation of the droplet
1374 interface, and its advection governs shape deformation. The final equation enforces curvature-related
1375 surface force balance, where σ is the surface tension coefficient, κ is the local curvature, and γ_{surface}
1376 represents surface forcing.1377 The spatial domain is discretized using a structured mesh aligned with the tokamak cross-section.
1378 A schematic of the cross-sectional mesh layout is shown in Figure C.2. The state space includes
1379 grid-sampled values of ρ , \mathbf{v} , \mathbf{B} , ϕ , and other derived quantities such as vorticity and magnetic energy.
1380 Control is applied via a discrete set of actions adjusting time-varying coil currents $\mathbf{I}_{\text{coil}}(t)$, which
1381 shape the confinement field. At each decision step, the reinforcement learning agent selects a control
1382 signal that reshapes the magnetic confinement field.1383 The reward function is defined as the negative of the minimum distance between the plasma droplet
1384 and the tokamak boundary, promoting stable suspension and avoiding wall contact. The long-term
1385 objective is to maintain the droplet within the desired confinement region while suppressing shape
1386 instabilities and maximizing equilibrium duration.

1387

1401 Figure C.2: Structured mesh layout of the tokamak poloidal cross-section.
1402
1403

1404 C.2.2 ERROR COUPLING BETWEEN RL AND PDE SPACES IN AI4S SYSTEMS
1405

1406 In this section, we analyze the sources and scales of error introduced in AI4S systems, especially the
1407 bidirectional projection between the RL space and the PDE-based physical space. As discussed in the
1408 main text and lemma S.3, for systems with observation noise, achieving an unbiased estimation of the
1409 RL transition kernel imposes a lower bound on the temporal resolution of the RL system, which is
1410 determined by the intrinsic dynamics of the underlying physical process. Based on this property, we
1411 couple the errors in the RL space and the physical space, and ultimately derive an expression for the
1412 total system error. Then through a $\rho - K$ analysis, we reveal the key factors that govern the overall
1413 numerical error of the AI4S system.

1414 We begin by introducing key notations to distinguish between two discretization structures:

1415 The **RL space** is characterized by (H_r, S_r, A_r) , denoting the RL decision horizon, state discretization,
1416 and action granularity. The **physical (PDE) space** is represented by (H_w, S_w, A_w) , corresponding to
1417 the physical evolution horizon, spatial resolution, and control parameter resolution in the underlying
1418 PDE model.

1419 Without loss of generality, we assume that the spatial grid resolution satisfies $\Delta x = \Delta y$. Accordingly,
1420 we define $\Delta x_{p,\text{int}}$ as the spatial grid spacing inside the PDE-based surrogate environment, and $\Delta x_{p,\text{bd}}$
1421 as the spatial spacing along the domain boundary. Similarly, $\Delta x_{r,\text{int}}$ and $\Delta x_{r,\text{bd}}$ denote the interior
1422 and boundary grid resolutions within the RL environment, respectively.

1423 Let the reinforcement learning policy operate on a coarse temporal scale, issuing actions every Δt_r
1424 seconds, while the physical system evolves at a much finer resolution $\Delta t_w \ll \Delta t_r$. Consequently,
1425 each RL control interval consists of $N_t = \Delta t_r / \Delta t_w$ internal PDE integration steps. The total
1426 prediction error associated with a single RL-AI4S step is then the cumulative result of local errors
1427 incurred at each fine-grained PDE time step.

1428 We now analyze the composition of prediction error in AI4S systems. During each RL interaction
1429 cycle, the observed state must be projected from the RL space to the PDE space, evolved forward
1430 under the PDE dynamics, and then projected back into the RL space. This bi-directional mapping
1431 introduces two distinct projection errors, both of which may be amplified near domain boundaries
1432 due to the nonlinear characteristics of the governing equations.

1433 Therefore, the total error consists of three main components: discretization error intrinsic to the RL
1434 space, numerical integration error accumulated within the PDE solver and coupling error caused by
1435 the nontrivial projections between RL and PDE spaces.

1437 **Fine-Grained Surrogate Error in PDE-Based Surrogate Environments.** We now analyze the
1438 surrogate error incurred at each fine-grained PDE step within the AI4S framework. This local error,
1439 denoted as $\Delta \phi_i^{(k)}$, arises from three primary sources: discretization in the spatial domain, imprecision
1440 in the control action, and interpolation near the physical boundary.

1442 We consider the level-set PDE that governs the evolution of an interface function ϕ : $\frac{\partial \phi}{\partial t} + \mathbf{v} \cdot \nabla \phi = 0$,
1443 where $\phi = \phi(x, t)$ is the level-set function and $\mathbf{v} = \mathbf{v}(x, t)$ is the velocity field. This PDE is
1444 discretized in time using a forward Euler method and in space using central differences.

1446 The numerical scheme used to evolve ϕ at a single time step is: $\frac{\phi_i^{n+1} - \phi_i^n}{\Delta t_w} + \mathbf{v}_i \cdot \nabla_h \phi_i^n = 0$, where
1447 $\nabla_h \phi_i^n$ is the discrete approximation of the gradient at grid point x_i and time t_n , computed using
1448 central differences.

1449 We now analyze the error introduced by this scheme, which is defined as the residual obtained by
1450 substituting the exact solution $\phi(x, t)$ into the numerical scheme.

1451 **Temporal truncation error:** Applying the Taylor expansion of $\phi(x, t)$ in time at point x_i gives:

$$1454 \phi_i^{n+1} = \phi_i^n + \Delta t_w \frac{\partial \phi}{\partial t} \Big|_i^n + \frac{\Delta t_w^2}{2} \frac{\partial^2 \phi}{\partial t^2} \Big|_i^n + \mathcal{O}(\Delta t_w^3).$$

1457 The finite difference approximation of the time derivative becomes:

1458

1459

1460

1461

1462

$$\frac{\phi_i^{n+1} - \phi_i^n}{\Delta t_w} = \frac{\partial \phi}{\partial t} \Big|_i^n + \frac{\Delta t_w}{2} \frac{\partial^2 \phi}{\partial t^2} \Big|_i^n + \mathcal{O}(\Delta t_w^2).$$

1463

To establish the temporal truncation error, we consider the second-order Taylor expansion of the level-set function $\phi(x, t)$ in time at a fixed spatial point x_i . In accordance with the physical model of tokamak plasma dynamics, the underlying equations exhibit parabolic-type behavior with smooth forcing and magnetic confinement fields. As such, we assume the solution $\phi(x, t)$ possesses sufficient temporal regularity.

1464

1465

1466

1467

In particular, we assume $\phi \in C^2([0, T] \times \Omega)$, which ensures that the second-order time derivative is uniformly bounded. That is, there exists a constant $C > 0$ such that: $\left| \frac{\partial^2 \phi}{\partial t^2}(x, t) \right| \leq C, \quad \forall (x, t) \in \Omega_T$. Under this regularity assumption, the finite difference approximation of the time derivative yields a local truncation error of the form: $\tau_{\text{time}} = \mathcal{O}(\Delta t_w)$, where the hidden constant depends on the supremum of $|\partial^2 \phi / \partial t^2|$.

1468

1469

1470

1471

1472

Spatial truncation error: We consider the approximation of the spatial gradient $\nabla \phi$ using central differences on a uniform grid with spacing $\Delta x = \Delta y = h$. For a function $\phi \in H^1(\Omega)$, we focus on the discrete gradient approximation defined as: $\frac{\partial_h \phi}{\partial x} \Big|_i := \frac{\phi(x_{i+1}) - \phi(x_{i-1})}{2h}$.

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

To evaluate the truncation error, we apply the first-order Taylor expansion $\phi(x_{i \pm 1}) = \phi(x_i) \pm h\phi'(x_i) + \mathcal{O}(h^2)$. Substituting into the difference formula yields $\frac{\phi(x_{i+1}) - \phi(x_{i-1})}{2h} = \phi'(x_i) + \mathcal{O}(h)$, which implies $\frac{\partial_h \phi}{\partial x} \Big|_i = \frac{\partial \phi}{\partial x} \Big|_i + \mathcal{O}(h)$. A similar result holds for the y -component. Combining both components, we obtain the discrete gradient $\nabla_h \phi = \nabla \phi + \mathcal{O}(h)$, which holds pointwise under the assumed regularity.

In the context of PDE-based surrogate environments, the spatial truncation error depends on the local grid region under consideration. Specifically, if we focus on the interior of the domain where the solution is sufficiently regular, the central difference approximation yields: $\nabla_h \phi = \nabla \phi + \mathcal{O}(\|\Delta \mathbf{x}_{p,\text{int}}\|)$. However, near the boundary or interface, the regularity of ϕ is reduced due to physical or geometric discontinuities. In such cases, the interpolation error is governed by the trace theorem, and satisfies only: $\|\phi - \phi_h\|_{L^2(\partial\Omega)} = \mathcal{O}(\|\Delta \mathbf{x}_{p,\text{bd}}\|^{1/2})$. For further details on the trace theorem and its implications for boundary error analysis, see Appendix D.

Action-induced control error: In our system, the velocity field \mathbf{v} , which governs the evolution of the plasma interface, is not directly actuated but indirectly induced through a cascade of physical interactions governed by the magnetohydrodynamic (MHD) equations. Specifically, the control input a_w corresponds to the current applied to external boundary coils. The discretization of the action space introduces finite-resolution perturbations Δa_w , which in turn cause quantization-induced errors in the boundary actuation.

The control-to-interface influence chain is as follows:

$$\Delta a_w \rightarrow \mathbf{B}|_{\partial\Omega} \rightarrow \mathbf{J} = \nabla \times \mathbf{B} \rightarrow \mathbf{v} \rightarrow \phi.$$

The applied control signal perturbs the boundary magnetic field $\mathbf{B}|_{\partial\Omega}$ through electromagnetic models such as the Biot–Savart law or magnetic vector potentials $\delta \mathbf{B}|_{\partial\Omega} = \mathcal{O}(\Delta a_w)$. The magnetic field \mathbf{B} evolves under the time-dependent induction equation $\frac{\partial \mathbf{B}}{\partial t} = \nabla \times (\mathbf{v} \times \mathbf{B}) + \eta \nabla^2 \mathbf{B}$. This is a nonlinear parabolic PDE, where the diffusion term $\eta \nabla^2 \mathbf{B}$ dominates the short-time behavior. Over a single control step, we can approximate the system as quasi-static and neglect the time derivative, resulting in $-\eta \nabla^2 \mathbf{B} \approx \nabla \times (\mathbf{v} \times \mathbf{B})$. Under this elliptic approximation, boundary disturbances introduced by Δa_w propagate into the interior via the smoothing effect of the Laplacian. By elliptic regularity and the trace theorem, we obtain: $\delta \mathbf{B}|_{\Omega} = \mathcal{O}(\Delta a_w^{1/2})$.

The momentum dynamics are governed by the MHD Navier–Stokes equation:

$$\rho \left(\frac{\partial \mathbf{v}}{\partial t} + \mathbf{v} \cdot \nabla \mathbf{v} \right) = -\nabla p + \mathbf{J} \times \mathbf{B} + \mu \nabla^2 \mathbf{v} + \mathbf{F}_{\text{surface}},$$

1512 where the Lorentz force $\mathbf{J} \times \mathbf{B}$ is the dominant actuator of plasma motion. Since $\mathbf{J} = \nabla \times \mathbf{B}$, both
 1513 $\delta\mathbf{B}$ and $\delta\mathbf{J}$ scale as $\mathcal{O}(\Delta a_w^{1/2})$. This yields a force perturbation:
 1514

$$1515 \quad \delta\mathbf{F}_{\text{Lorentz}} = \delta\mathbf{J} \times \mathbf{B} + \mathbf{J} \times \delta\mathbf{B} = \mathcal{O}(\Delta a_w^{1/2}),$$

1516 and correspondingly a sublinear response in velocity: $\Delta\mathbf{v} = \mathcal{O}(\Delta a_w^{1/2})$.
 1517

1518 **Final per-step surrogate error:** We now summarize the local surrogate error incurred in the
 1519 discretized interface evolution governed by $\frac{\partial\phi}{\partial t} + \mathbf{v} \cdot \nabla\phi = 0$. Each term introduces a distinct
 1520 approximation error: The time derivative $\partial_t\phi$ is discretized by forward Euler, incurring $\mathcal{O}(\Delta t_w)$
 1521 truncation error. The spatial gradient $\nabla\phi$ is computed via finite differences, contributing $\mathcal{O}(\|\Delta\mathbf{x}_{p,\text{int}}\|)$
 1522 in the interior and $\mathcal{O}(\|\Delta\mathbf{x}_{p,\text{bd}}\|^{1/2})$ near boundaries. The velocity field \mathbf{v} is perturbed by action
 1523 discretization, with induced error $\mathcal{O}(\Delta a_w^{1/2})$.
 1524

1525 Substituting these into the PDE gives:

$$1526 \quad \frac{\phi^{n+1} - \phi^n}{\Delta t_w} + (\mathbf{v} + \mathcal{O}(\Delta a_w^{1/2})) \cdot (\nabla\phi + \mathcal{O}(\|\Delta\mathbf{x}_{p,\text{bd}}\|^{1/2})),$$

1528 yielding the total local surrogate error per step as:
 1529

$$1530 \quad \Delta_{\phi}^{(k)} = \mathcal{O}(\Delta t_w) + \mathcal{O}(\|\Delta\mathbf{x}_{p,\text{int}}\|) + \mathcal{O}(\Delta a_w^{1/2}) + \mathcal{O}(\|\Delta\mathbf{x}_{p,\text{bd}}\|^{1/2}) + \mathcal{O}(\Delta a_w^{1/2} \cdot \|\Delta\mathbf{x}_{p,\text{bd}}\|^{1/2}).$$

1532 **One-Step Prediction Error in RL-Based Control Environments.** In the AI4S control framework,
 1533 the RL agent interacts with a PDE-based simulator by issuing a discretized action a_r based on an
 1534 observed state s_r , which itself is a downsampled or filtered version of the true PDE state $y \in \mathcal{Y}$. The
 1535 PDE system then evolves the dynamics over a time horizon Δt_r and returns a new state s'_r . The
 1536 surrogate prediction error in this interaction arises from discretization of state and action inputs and
 1537 integration-induced temporal accumulation.

1538 **RL observation error as truncation error in PDE-Based environments:** In PDE-based reinforcement
 1539 learning environments, the agent does not observe the full continuous physical state $y(x) \in \mathbb{R}^d$,
 1540 but rather a discretized version $s_r \in \mathbb{R}^d$, obtained via sampling or interpolation over a coarse spatial
 1541 mesh. This introduces a spatial offset between the actual physical location x and the evaluation point
 1542 \tilde{x} used in downstream PDE computation: $\|x - \tilde{x}\| \leq \|\Delta\mathbf{x}_{r,\text{int}}\|$. Such an offset mimics the numerical
 1543 behavior of spatial truncation errors in classical finite difference schemes.

1544 Consider a PDE used for control computation in the simulator: $\frac{\partial\phi}{\partial t} + \mathbf{v} \cdot \nabla\phi = 0$, the derivative $\nabla\phi$ is
 1545 evaluated not at the true point x , but at \tilde{x} . The following analysis parallels our earlier fine-grained
 1546 surrogate error in PDE-based surrogate environments, from which we conclude that the truncation
 1547 error induced by RL's observation of the state scales as $\mathcal{O}(\|\Delta\mathbf{x}_{r,\text{int}}\|)$.
 1548

1549 Similarly, near the domain boundary $\partial\Omega$, the solution $\phi \in H^1(\Omega)$ typically lacks full smoothness.
 1550 By the trace theorem, the restriction of ϕ to the boundary lies in the fractional Sobolev space: $\phi|_{\partial\Omega} \in$
 1551 $H^{1/2}(\partial\Omega)$. Consequently, interpolation errors on the boundary are governed by lower regularity,
 1552 yielding $\|\phi - \phi_h\|_{L^2(\partial\Omega)} = \mathcal{O}(\|\Delta\mathbf{x}_{r,\text{bd}}\|^{1/2})$, which mirrors the $\mathcal{O}(h^{1/2})$ rate of convergence found
 1553 in boundary discretization analysis of finite difference schemes.

1554 In summary, the state observation errors in RL, when incorporated into the PDE solver, act as implicit
 1555 spatial discretization errors. Their scaling matches classical truncation theory:
 1556

$$1557 \quad \Delta_{\text{state}} = \begin{cases} \mathcal{O}(\|\Delta\mathbf{x}_{r,\text{int}}\|), & \text{for interior evaluation,} \\ \mathcal{O}(\|\Delta\mathbf{x}_{r,\text{bd}}\|^{1/2}), & \text{for boundary evaluation.} \end{cases}$$

1559 **Action discretization error:** The agent selects an action a_r from a discretized control space, which
 1560 governs the boundary actuation of the PDE system. The true control variable a is approximated by a_r ,
 1561 with finite resolution Δa_r . This discretization introduces quantization error in the applied boundary
 1562 signal.

1564 Similar to our earlier fine-grained surrogate error analysis in PDE-based environments, we observe
 1565 that this control mismatch propagates through the MHD system in a nonlinear and spatially diffused
 1566 manner. The resulting perturbation in the velocity field—responsible for evolving the system

state—scales sublinearly with respect to the action resolution: $\Delta a_r \Rightarrow \delta v = \mathcal{O}(\Delta a_r^{1/2})$. This velocity perturbation directly enters the PDE dynamics and contributes to the overall one-step prediction error in the RL loop.

Temporal integration error: During each RL step of duration Δt_r , the PDE simulator executes $N_t = \Delta t_r / \Delta t_w$ fine-grained steps with timestep Δt_w . At each substep, the forward Euler integration induces local truncation error: $\tau_t^{(n)} = \mathcal{O}(\Delta t_w)$, so that the cumulative integration error over one RL step becomes: $\sum_{n=1}^{N_t} \tau_t^{(n)} = \mathcal{O}(\Delta t_r)$.

Summing all contributions, the one-step surrogate prediction error in RL space satisfies:

$$\Delta_{\text{RL}} = \underbrace{C_1 \|\Delta \mathbf{x}_{r,\text{int}}\|}_{\text{State (interior)}} + \underbrace{C_2 \|\Delta \mathbf{x}_{r,\text{bd}}\|^{1/2}}_{\text{State (boundary)}} + \underbrace{C_3 \Delta a_r^{1/2}}_{\text{Action resolution}} + \underbrace{C_4 \Delta t_r}_{\text{Temporal propagation}}.$$

Refined Total Error Decomposition with Time-Scale Separation. Based on our previous analysis of surrogate error in both the PDE and RL components, we now integrate all sources of discretization and numerical error into a unified framework. Over a single RL step, the total prediction error can be decomposed as follows:

$$\begin{aligned} \Delta_{\text{total}} = & \underbrace{C_1 \|\Delta \mathbf{x}_{r,\text{int}}\|}_{\text{RL (interior space)}} + \underbrace{C_2 \|\Delta \mathbf{x}_{r,\text{bd}}\|^{1/2}}_{\text{RL (boundary space)}} + \underbrace{C_3 \Delta a_r^{1/2}}_{\text{RL (action space)}} + \underbrace{C_4 \Delta t_r}_{\text{RL (time)}} \\ & + \underbrace{C_5 \frac{\Delta t_r}{\Delta t_w} \left(\|\Delta \mathbf{x}_{p,\text{int}}\| + \Delta a_w^{1/2} + \|\Delta \mathbf{x}_{p,\text{bd}}\|^{1/2} + \Delta a_w^{1/2} \cdot \|\Delta \mathbf{x}_{p,\text{bd}}\|^{1/2} \right)}_{\text{PDE surrogate error (space + action + boundary), amplified by time scale separation}}. \end{aligned} \quad (\text{C.6})$$

To ensure that the learned RL transition kernel remains robust under observation uncertainty Δy , we require that: $\Delta_{\text{total}} = \mathcal{O}(\Delta y)$. This constraint implies a resolution matching condition across all discretization dimensions. Specifically, we obtain the following asymptotic scaling relations:

$$\begin{aligned} \|\Delta \mathbf{x}_{r,\text{int}}\| &\sim \|\Delta \mathbf{x}_{p,\text{int}}\| \sim \Delta a_r^{1/2} \sim \Delta a_w^{1/2} \sim \Delta t_r \sim \Delta y, \\ \|\Delta \mathbf{x}_{r,\text{bd}}\| &\sim \|\Delta \mathbf{x}_{p,\text{bd}}\| \sim \Delta y^2, \quad \Delta t_w \ll \Delta y. \end{aligned} \quad (\text{C.7})$$

ρ -K Analysis. We now turn to the analysis of the error rate associated with predicting the next state in AI4S environments based on an RL agent’s current observation and action. Specifically, we define a prediction error event as one where the next predicted state fails to fall within the correct discretized RL state cell. Accordingly, we define the relative prediction error rate as: $\rho = 1 - \frac{\Delta y}{\text{Total Prediction Error}}$, where Δy is the size of the RL spatial grid. In a three-dimensional state space, it corresponds to the cube of the grid length.

The total prediction error in the system arises from a combination of three interrelated sources. First, observation noise introduces inherent uncertainty, stemming from either partial observability or sensor imprecision within the reinforcement learning environment. Second, numerical surrogate error which quantified earlier in Equation equation C.6 captures the discretization-induced inaccuracies originating from both the RL and PDE components. Finally, the intrinsic growth of initial perturbations, as detailed in Lemma S.3, contributes an additional error term due to the modal amplification behavior of the MHD system; specifically, this leads to a forward-propagated uncertainty of order $\lambda_1 \Delta y / H_r$, where λ_1 denotes the leading eigenvalue and H_r is the RL planning horizon.

Now we further quantify the numerical surrogate error based on the resolution matching relations previously derived. Given the grid refinement ratio $K := \Delta \mathbf{x}_r / \Delta \mathbf{x}_p$, which compares the spatial resolution of the PDE environment to that of the RL environment, and considering the combined error impact across three dimensions in state space, we obtain: $\Delta_{\text{total}} \sim \frac{C_1}{K^3} \Delta y^{(j)}$. Substituting all three error components into the definition of ρ , we obtain:

$$\rho = 1 - \frac{\Delta y}{\lambda_1 \Delta y / H_r + \Delta y + \Delta_{\text{total}}} = 1 - \frac{1}{\lambda_1 / H_r + 1 + C_1 / K^3}.$$

So in the limit of high-resolution settings, i.e., when H_r and K are large, we obtain $\rho = \mathcal{O}\left(\frac{1}{H_r} + \frac{1}{K^3}\right)$.

1620
1621

C.2.3 OPTIMAL COMPUTATIONAL COST ALLOCATION BETWEEN RL AND AI4S

1622
1623
1624

In this section, building on the previous analysis of error scaling, we provide a detailed explanation of how to align the discretization parameters of reinforcement learning and AI4S models to achieve optimal computational efficiency.

1625
1626
1627
1628
1629

It is evident that higher fidelity in the physical model leads to more accurate predictions, which in turn accelerates convergence to the optimal policy in RL. However, since AI4S simulators are typically executed under fixed computational budgets, it is crucial to ensure that the computational load in the RL loop is of the same order as that of the AI4S module. Our objective is to minimize total computational cost by optimally distributing resolution between the two systems.

1630
1631
1632
1633
1634

According to Lemma B.1, the sample complexity or computational cost of the RL component can be expressed as $\text{Cost}_{\text{RL}} = \mathcal{O}\left(\frac{H_r^3 S_r A_r \cdot \log(1/\delta)}{\varepsilon^2}\right)$. Meanwhile, the computational cost of the AI4S physical simulator is governed by its spatiotemporal resolution, and can be approximated by: $\text{Cost}_{\text{AI4S}} = \mathcal{O}(H_w S_w A_w)$.

1635
1636
1637
1638

To ensure computational consistency, we aim to align these quantities. This alignment is further guided by the theoretical result in Lemma S.3, where the RL horizon H_r is constrained by the intrinsic dynamics of the PDE system: $H_r \lesssim \frac{1}{\lambda_1}$, where λ_1 is the dominant growth rate in the linearized MHD system.

1639
1640

Based on the resolution matching conditions derived earlier, we assume the following scalings:

1641
1642

$$A_r = \mathcal{O}(A_w) = \mathcal{O}(x_w^2), S_r = \mathcal{O}(S_w) = \mathcal{O}(x_w^6), H_r = H_w^{1/3} \cdot \mathcal{O}(x_w).$$

1643
1644
1645
1646
1647

Here, x_w denotes the effective resolution of the AI4S spatial grid. The scaling $S_r = \mathcal{O}(x_w^6)$ arises from both the intrinsic dimensionality of the state space and the additional resolution requirements near the domain boundary. In tokamak plasma systems, the physical state space is inherently three-dimensional. Under uniform discretization, this implies that the number of internal state variables scales as $\mathcal{O}(x_w^3)$.

1648
1649
1650
1651
1652
1653

Besides, near the boundary $\partial\Omega$, the PDE solution often exhibits reduced regularity—formally, if $y \in H^1(\Omega)$, then its restriction to the boundary lies in the fractional Sobolev space $H^{1/2}(\partial\Omega)$, as stated by the trace theorem. This lower regularity implies that to maintain the same level of approximation accuracy at the boundary, the spatial grid must be refined further. In particular, the resolution must be squared in order to compensate for the smoothness loss, effectively contributing an additional factor of $\mathcal{O}(x_w^3)$ from the boundary region.

1654
1655
1656
1657
1658

The total RL cost then becomes:

$$\text{Cost}_{\text{RL}} \propto \frac{H_r S_r A_r \cdot \log(1/\delta)}{(\Delta y^{(j)} - p_{\max}^{(j)})^2} \approx \frac{H_r^9 K^8}{\varepsilon} \cdot \left(\frac{1}{1 - \frac{1}{H_r} - \frac{1}{K^3}} \right)^2, \quad (\text{C.8})$$

1659

Our goal is to find the optimal refinement ratio K^* that minimizes this cost expression. Formally, we solve the following optimization problem:

1660
1661
1662
1663
1664

$$\min_{K>0} C(K) := \frac{H_r^9 K^8}{\varepsilon} \cdot \left(\frac{1}{1 - \frac{1}{H_r} - \frac{1}{K^3}} \right)^2.$$

1665
1666
1667
1668
1669
1670
1671

To find the stationary point, we differentiate $C(K)$ with respect to K and set the derivative to zero: $\frac{dC}{dK} = 0$. Denote $\alpha := 1 - \frac{1}{H_r}$, then $C(K) \propto K^8 \left(\frac{1}{\alpha - \frac{1}{K^3}} \right)^2$. Taking logarithmic derivative: $\frac{d \log C}{dK} = \frac{8}{K} - 2 \cdot \frac{3}{K^4(\alpha - \frac{1}{K^3})}$. Setting this derivative to zero and solving yields the optimal K^* : $K^* = \left(\frac{7}{4(1 - \frac{1}{H_r})} \right)^{1/3}$. So in the high-resolution regime where $H_r \gg 1$, we approximate:

1672
1673

$$K^* \gtrsim \left(\frac{7}{4(1 - \frac{1}{\lambda_1})} \right)^{1/3} \approx \left(\frac{7}{4} e^{1/\lambda_1} \right)^{1/3}.$$

1674 C.3 TURBULENT AIRFOIL REGULATION
16751676 C.3.1 ENVIRONMENT DESCRIPTION
1677

1678 In this task, we study an unsteady aerodynamic system involving the control of a turbulent airfoil
1679 flow field. The objective is to regulate the pitching angle of the airfoil to manipulate the surrounding
1680 flow structure and ultimately maximize lift (Portal-Porras et al., 2022).

1681 The governing equations for the fluid dynamics are the two-dimensional incompressible Navier-Stokes
1682 equations:

1683

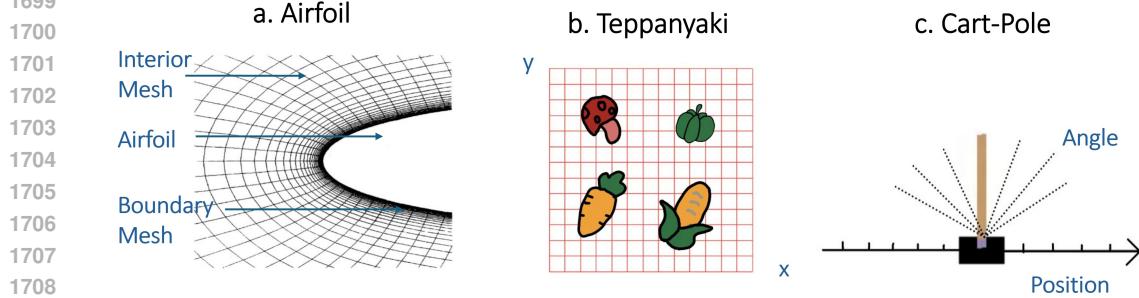
$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \nu \Delta \mathbf{u}$$

1684 where $\mathbf{u}(x, y, t) = (u(x, y, t), v(x, y, t))$ is the velocity vector field, $p(x, y, t)$ denotes the pressure, ν
1685 is the kinematic viscosity, ∇ is the gradient operator, and Δ is the Laplacian operator.

1686 The system is also subject to the incompressibility constraint: $\nabla \cdot \mathbf{u} = 0$. These equations define the
1687 error sources and discretization structure that form the basis for our multi-resolution analysis.

1688 In this environment, we adopt a structured grid discretization for the spatial domain, employing an
1689 O-grid topology that conforms to the geometry of the airfoil. As illustrated in Figure C.3.a, the mesh
1690 is progressively refined near the surface of the airfoil to accurately resolve the boundary layer. The
1691 state space is thus defined over this discretized grid. The action space is discretized into a finite set
1692 of pitch angle increments. At each time step, the RL agent selects a discrete action corresponding
1693 to a rotation of the airfoil about its chord line. The reward function is constructed based on the
1694 instantaneous lift coefficient, with the overall objective of maximizing its long-term average.

1695



1711 Figure C.3: Visualization of Discretization Structures in Different Control Environments.

1712

1713 C.3.2 ERROR COUPLING BETWEEN RL AND PDE SPACES IN AI4S SYSTEMS
1714

1715 Since the state space of the turbulent airfoil flow field also features a division between interior and
1716 boundary regions, we adopt an analysis framework analogous to that used for the tokamak device to
1717 analyze the composition of prediction error in AI4S systems.

1718

1719 **Fine-Grained Surrogate Error in Airfoil Flow Environments.** We now analyze the surrogate
1720 error incurred at each fine-grained PDE step within AI4S framework. In airfoil flow control tasks, the
1721 dynamics are governed by the two-dimensional incompressible Navier-Stokes equations: $\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \nu \Delta \mathbf{u}$. This PDE is discretized in time using a forward Euler method and in space
1722 using central differences.

1723

1724 **Temporal truncation error:** Using the Taylor expansion at time t_n , we write:

1725

$$\mathbf{u}^{n+1} = \mathbf{u}^n + \Delta t_w \frac{\partial \mathbf{u}}{\partial t} \Big|_n + \frac{1}{2} \Delta t_w^2 \frac{\partial^2 \mathbf{u}}{\partial t^2} \Big|_n + \mathcal{O}(\Delta t_w^3).$$

1726

1727

1728 Then the numerical derivative satisfies:
 1729

$$1730 \quad \frac{\mathbf{u}^{n+1} - \mathbf{u}^n}{\Delta t_w} = \frac{\partial \mathbf{u}}{\partial t} \Big|_t^n + \frac{1}{2} \Delta t_w \frac{\partial^2 \mathbf{u}}{\partial t^2} \Big|_t^n + \mathcal{O}(\Delta t_w^2),$$

1732 which implies a local truncation error of $\mathcal{O}(\Delta t_w)$ provided $\mathbf{u} \in C^2$.
 1733

1734 **Spatial truncation error:** For a function $u \in H^1(\Omega)$, we focus on the numerical approximation
 1735 of the convective term $(\mathbf{u} \cdot \nabla) \mathbf{u}$ and the diffusive term $\nu \Delta \mathbf{u}$, which appear in the Navier-Stokes
 1736 momentum equation: $\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \nu \Delta \mathbf{u}$. For the convective term, take one component
 1737 $u(x, t)$, and use central differences to approximate: $u \frac{\partial u}{\partial x} \Big|_{x_i} \approx u(x_i) \cdot \frac{u(x_{i+1}) - u(x_{i-1})}{2h}$. Since the
 1738 function $u \in H^1(\Omega)$, the accuracy of the Taylor expansion is limited to $\mathcal{O}(h^2)$. Applying Taylor
 1739 expansion, we obtain: $\frac{u(x_{i+1}) - u(x_{i-1})}{2h} = u'(x_i) + \mathcal{O}(h)$, which leads to: $u(x_i) \cdot \frac{u(x_{i+1}) - u(x_{i-1})}{2h} =$
 1740 $u(x_i)u'(x_i) + \mathcal{O}(h)$.
 1741

1742 For the viscous diffusion term, the Laplacian Δu is approximated using second-order central
 1743 differences: $\frac{\partial^2 u}{\partial x^2} \Big|_{x_i} \approx \frac{u(x_{i+1}) - 2u(x_i) + u(x_{i-1})}{h^2}$, and similarly in the y -direction. From Taylor expansion:
 1744 $\frac{u(x_{i+1}) - 2u(x_i) + u(x_{i-1})}{h^2} = u''(x_i) + \mathcal{O}(h)$. Hence, the Laplacian approximation also introduces
 1745 an error of $\mathcal{O}(h)$.
 1746

1747 Therefore, both nonlinear convective and diffusive terms yield local discretization errors that can be
 1748 bounded as:
 1749

$$1750 \quad (\mathbf{u} \cdot \nabla_h) \mathbf{u} = (\mathbf{u} \cdot \nabla) \mathbf{u} + \mathcal{O}(\|\Delta \mathbf{x}_{p,\text{int}}\|), \quad \Delta_h \mathbf{u} = \Delta \mathbf{u} + \mathcal{O}(\|\Delta \mathbf{x}_{p,\text{int}}\|).$$

1752 On or near the airfoil surface $\partial\Omega$, the velocity field \mathbf{u} may possess only H^1 regularity. By the trace
 1753 theorem, its restriction belongs to: $\mathbf{u}|_{\partial\Omega} \in H^{1/2}(\partial\Omega)$, and the numerical interpolation error along
 1754 the boundary grid satisfies: $\|\mathbf{u} - \mathbf{u}_h\|_{L^2(\partial\Omega)} = \mathcal{O}(\|\Delta \mathbf{x}_{p,\text{bd}}\|^{1/2})$.
 1755

1756 **Action-induced control error:** In airfoil flow control, boundary actuation modifies wall velocities.
 1757 Suppose this control is discretized with resolution Δa_w . The boundary perturbation in velocity is:
 1758 $\delta \mathbf{u}|_{\partial\Omega} = \mathcal{O}(\Delta a_w)$. Because the viscous diffusion equation governs short-time dynamics: $\frac{\partial \mathbf{u}}{\partial t} =$
 1759 $\nu \Delta \mathbf{u}$, we model the induced field internally by elliptic smoothing, yielding: $\delta \mathbf{u}|_{\Omega} = \mathcal{O}(\Delta a_w^{1/2})$.
 1760

1761 **Final surrogate error per step:** Combining the above terms, the total local surrogate error per PDE
 1762 step in the airfoil flow simulator is:
 1763

$$1763 \quad \Delta_{\mathbf{u}}^{(k)} = \mathcal{O}(\Delta t_w) + \mathcal{O}(\|\Delta \mathbf{x}_{p,\text{int}}\|) + \mathcal{O}(\Delta a_w^{1/2}) + \mathcal{O}(\|\Delta \mathbf{x}_{p,\text{bd}}\|^{1/2}).$$

1764 **One-Step Prediction Error in RL-Controlled Airfoil Environments.** Based on our earlier sur-
 1765rogate error decomposition for tokamak systems, we now analyze the one-step prediction error
 1766 in RL-based control of unsteady airfoil flow. The underlying dynamics are governed by the 2D
 1767 incompressible Navier-Stokes equations: $\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \nu \Delta \mathbf{u}$. In this setting, the RL agent
 1768 selects an action a_r based on an observed state s_r , which is a spatially discretized and possibly noisy
 1769 version of the continuous flow field \mathbf{u} . The PDE simulator uses a_r to modify boundary conditions,
 1770 evolves the fluid state over a macro timestep Δt_r , and returns the next observed state s'_r .
 1771

1772 **RL observation error as truncation error:** Due to spatial discretization, the observation s_r deviates
 1773 from the true field $\mathbf{u}(x)$ by at most the grid resolution: $\|x - \tilde{x}\| \leq \|\Delta \mathbf{x}_{r,\text{int}}\|$. This offset behaves
 1774 analogously to a spatial truncation error in finite difference schemes. As previously analyzed, for
 1775 smooth internal flow: $\Delta_{\text{state}}^{\text{int}} = \mathcal{O}(\|\Delta \mathbf{x}_{r,\text{int}}\|)$. Near the airfoil surface, where boundary layers reduce
 1776 regularity, trace theorem implies: $\|\mathbf{u} - \mathbf{u}_h\|_{L^2(\partial\Omega)} = \mathcal{O}(\|\Delta \mathbf{x}_{r,\text{bd}}\|^{1/2})$.
 1777

1778 **Action discretization error:** The discretized control signal a_r influences boundary actuation. Finite
 1779 resolution Δa_r introduces quantization error, which affects boundary velocity: $\delta \mathbf{u}|_{\partial\Omega} = \mathcal{O}(\Delta a_r)$.
 1780 Via the viscous smoothing governed by $\nu \Delta \mathbf{u}$, this error diffuses into the interior: $\delta \mathbf{u}|_{\Omega} = \mathcal{O}(\Delta a_r^{1/2})$.
 1781

1782 **Temporal integration error:** Each RL step of length Δt_r consists of $N_t = \Delta t_r / \Delta t_w$ fine-scale
 1783 Euler updates with local error $\mathcal{O}(\Delta t_w)$. Total accumulation yields: $\sum_{n=1}^{N_t} \tau_t^{(n)} = \mathcal{O}(\Delta t_r)$.
 1784

1782 **Final surrogate error:** Combining all contributions, we obtain the total RL one-step prediction
 1783 error:

$$1784 \Delta_{\text{RL}} = \underbrace{C_1 \|\Delta \mathbf{x}_{r,\text{int}}\|}_{\text{State (interior)}} + \underbrace{C_2 \|\Delta \mathbf{x}_{r,\text{bd}}\|^{1/2}}_{\text{State (boundary)}} + \underbrace{C_3 \Delta a_r^{1/2}}_{\text{Action resolution}} + \underbrace{C_4 \Delta t_r}_{\text{Temporal integration}}.$$

$$1785$$

$$1786$$

1787 **Refined Total Error Decomposition with Time-Scale Separation.** Based on our prior analysis of
 1788 surrogate errors in both PDE and RL components, we now present a unified decomposition of all
 1789 numerical and discretization errors. Over a single RL step, the total prediction error can be expressed
 1790 as:

$$1791$$

$$1792 \Delta_{\text{total}} = \underbrace{C_1 \|\Delta \mathbf{x}_{r,\text{int}}\|}_{\text{RL (interior space)}} + \underbrace{C_2 \|\Delta \mathbf{x}_{r,\text{bd}}\|^{1/2}}_{\text{RL (boundary space)}} + \underbrace{C_3 \Delta a_r^{1/2}}_{\text{RL (action space)}} + \underbrace{C_4 \Delta t_r}_{\text{RL (time)}} \\ 1793 \\ 1794 \\ 1795 \\ 1796 \\ 1797 \\ 1798 \\ 1799 \\ 1800 \\ 1801 \\ 1802 \\ 1803 \\ 1804 \\ 1805 \\ 1806 \\ 1807 \\ 1808 \\ 1809 \\ 1810 \\ 1811 \\ 1812 \\ 1813 \\ 1814 \\ 1815 \\ 1816 \\ 1817 \\ 1818 \\ 1819 \\ 1820 \\ 1821 \\ 1822 \\ 1823 \\ 1824 \\ 1825 \\ 1826 \\ 1827 \\ 1828 \\ 1829 \\ 1830 \\ 1831 \\ 1832 \\ 1833 \\ 1834 \\ 1835$$

$$+ \underbrace{C_5 \frac{\Delta t_r}{\Delta t_w} \left(\|\Delta \mathbf{x}_{p,\text{int}}\| + \Delta a_w^{1/2} + \|\Delta \mathbf{x}_{p,\text{bd}}\|^{1/2} \right)}_{\text{PDE surrogate error amplified by time-scale separation}}. \quad (\text{C.9})$$

To maintain robustness of the RL transition kernel under observation uncertainty Δy , we impose the constraint: $\Delta_{\text{total}} = \mathcal{O}(\Delta y)$. This leads to the following resolution matching conditions across discretization components:

$$\|\Delta \mathbf{x}_{r,\text{int}}\| \sim \|\Delta \mathbf{x}_{p,\text{int}}\| \sim \Delta a_r^{1/2} \sim \Delta a_w^{1/2} \sim \Delta t_r \sim \Delta y, \\ \|\Delta \mathbf{x}_{r,\text{bd}}\| \sim \|\Delta \mathbf{x}_{p,\text{bd}}\| \sim \Delta y^2, \quad \Delta t_w \ll \Delta y. \quad (\text{C.10})$$

ρ -K Analysis. We now turn to the analysis of the error rate associated with predicting the next state in AI4S environments based on an RL agent’s current observation and action. We define the relative prediction error rate as: $\rho = 1 - \frac{\Delta y}{\text{Total Prediction Error}}$, where Δy is the size of the RL spatial grid. In a two-dimensional state space, it corresponds to the square of the grid length. Similar to the previous analysis, the total prediction error comprises three key components: observation noise, numerical surrogate error, and intrinsic growth of initial perturbation.

Now we further quantify the numerical surrogate error based on the resolution matching relations previously derived. Considering the combined error impact across two dimensions in state space, we obtain: $\Delta_{\text{total}} \sim \frac{C_1}{K^2} \Delta y^{(j)}$. Substituting all three error components into the definition of ρ , we obtain:

$$\rho = 1 - \frac{\Delta y}{\lambda_1 \Delta y / H_r + \Delta y + \Delta_{\text{total}}} = 1 - \frac{1}{\lambda_1 / H_r + 1 + C_1 / K^2}.$$

So in the limit of high-resolution settings, i.e., when H_r and K are large, we obtain: $\rho = \mathcal{O}\left(\frac{1}{H_r} + \frac{1}{K^2}\right)$. The subsequent results on optimal computational cost allocation are summarized in Table 1 of the main text. As this involves only straightforward calculations and scaling arguments, we omit detailed elaboration here.

1822 C.4 TEPPANYAKI HEAT SEQUENCING

1823 C.4.1 ENVIRONMENT DESCRIPTION

1826 In this task, we study a two-dimensional heat transfer system representing a teppanyaki cooking
 1827 surface. The goal is to use reinforcement learning to control the heat distribution across the iron plate
 1828 to ensure the fastest and most uniform cooking of food.

1829 The thermal evolution of the plate is governed by the two-dimensional heat diffusion equation:

$$1831 \frac{\partial T(x, y, t)}{\partial t} = \alpha \left(\frac{\partial^2 T}{\partial x^2} + \frac{\partial^2 T}{\partial y^2} \right) + Q(x, y, t),$$

$$1832$$

$$1833$$

1834 where $T(x, y, t)$ denotes the temperature at spatial position (x, y) and time t ; α is the thermal
 1835 diffusivity coefficient, determined by the material of the iron plate; $Q(x, y, t)$ represents the external
 heat source, which can be actively controlled by the agent at each grid cell (Koehler et al., 2024).

1836 The spatial domain is discretized using a uniform Cartesian grid. Each grid cell corresponds to
 1837 a controllable heating unit, allowing the RL agent to adjust the local heat input. As shown in
 1838 Figure C.3.b. The action space is discretized into finite heating levels for each control unit. The
 1839 overall objective is to optimize heating strategies that ensure efficient thermal convergence within the
 1840 shortest possible time.

1841

1842 C.4.2 ERROR COUPLING BETWEEN RL AND PDE SPACES IN AI4S SYSTEMS

1843

1844 Unlike the previously introduced two environments, the teppanyaki setting adopts a uniform Cartesian
 1845 grid discretization for the spatial domain. As a result, there is no distinction between interior and
 1846 boundary regions, and no precision loss due to boundary regularity degradation. In the following,
 1847 we adopt an analysis framework analogous to that used for the tokamak device to examine the
 1848 composition of prediction error in AI4S systems.

1849

1850 **Fine-Grained Surrogate Error in Teppanyaki Heat Diffusion Environment.** We now analyze
 1851 the local surrogate error incurred in each fine-grained PDE step under the teppanyaki heat diffusion
 1852 setting. The governing equation for the temperature field $T(x, y, t)$ over the iron plate is the two-
 1853 dimensional heat diffusion equation: $\frac{\partial T}{\partial t} = \alpha \Delta T + Q(x, y, t)$, where α is the thermal diffusivity and
 1854 Q is the external heating source, modulated by the control input.

1855

1856 **Temporal truncation error:** We discretize the temporal derivative using the forward Euler scheme:
 1857 $\frac{T_{i,j}^{n+1} - T_{i,j}^n}{\Delta t_w} = \alpha \Delta_h T_{i,j}^n + Q_{i,j}^n$, where $\Delta_h T_{i,j}^n$ is the discrete Laplacian at grid point i, j . Applying
 1858 Taylor expansion, the numerical derivative satisfies:

1859

$$\frac{\mathbf{T}_{i,j}^{n+1} - \mathbf{T}_{i,j}^n}{\Delta t_w} = \frac{\partial \mathbf{T}}{\partial t} \Big|_{i,j}^n + \frac{1}{2} \Delta t_w \frac{\partial^2 \mathbf{T}}{\partial t^2} \Big|_{i,j}^n + \mathcal{O}(\Delta t_w^2),$$

1860

1861 which implies a local truncation error of $\mathcal{O}(\Delta t_w)$ provided $\mathbf{T} \in C^2$. We obtain the temporal
 1862 truncation error: $\tau_{\text{time}} = \mathcal{O}(\Delta t_w)$.

1863

1864 **Spatial truncation error:** The Laplacian operator is approximated by central differences:

1865

$$\Delta T \approx \frac{T_{i+1,j} - 2T_{i,j} + T_{i-1,j}}{(\Delta x)^2} + \frac{T_{i,j+1} - 2T_{i,j} + T_{i,j-1}}{(\Delta y)^2}.$$

1866

1867 Assuming $T \in H^1(\Omega)$, the second-order spatial derivatives exist in the weak sense, but due to
 1868 limited regularity we expect that: $\Delta_h T = \Delta T + \mathcal{O}(\|\Delta \mathbf{x}_p\|)$. This yields a spatial truncation error:
 1869 $\tau_{\text{space}} = \mathcal{O}(\|\Delta \mathbf{x}_p\|)$.

1870

1871 **Action-induced control error:** In this environment, control is implemented by adjusting the external
 1872 heat source term $Q(x, y, t)$ in the heat diffusion equation. Let the continuous control signal be
 1873 denoted by $a(t)$, and suppose it is discretized with resolution Δa_w into a piecewise constant control
 1874 $a_w(t)$. Assuming that the mapping from the control input to the heat source is linear, we obtain:
 1875 $Q(x, y, t) = \bar{Q}(x, y, t) + \delta Q(x, y, t)$, where \bar{Q} is the nominal source corresponding to $a_w(t)$, and
 1876 the perturbation in the source is: $\delta Q(x, y, t) = f(\delta a) = \mathcal{O}(\Delta a_w)$.

1877

1878 Since the heat equation is linear in Q , the perturbation in Q propagates linearly to the temperature field.
 1879 Hence, the resulting temperature deviation induced by control discretization satisfies: $\delta T = \mathcal{O}(\Delta a_w)$.

1880

1881 **Total per-step surrogate error.** Combining all error sources, the total surrogate error in the tempera-
 1882 ture prediction per PDE step is: $\Delta_T^{(k)} = \mathcal{O}(\Delta t_w) + \mathcal{O}(\|\Delta \mathbf{x}_p\|) + \mathcal{O}(\Delta a_w)$.

1883

1884

1885

1886

1887

1888 **One-Step Prediction Error in RL-Controlled Heat Diffusion Environments.** In the teppanyaki
 1889 plate task modeled by the heat diffusion equation, the RL agent interacts with a PDE simulator by
 1890 issuing a discretized action a_r and observing a discretized temperature field s_r , which is derived from
 1891 the true temperature field $T(x, y, t)$ through uniform spatial sampling. The simulator then evolves
 1892 the system over a time step Δt_r and returns an updated observation s'_r . The one-step prediction error
 1893 in this process arises from three main sources: state observation error, action discretization error, and
 1894 temporal integration error.

1895

1896

1897

1898 **RL observation error:** The observed state s_r is obtained by projecting the continuous temperature
 1899 field $T(x, y, t)$ onto a coarsely discretized uniform Cartesian grid. This introduces a spatial obser-
 1900 vation error due to the offset between the true location (x, y) and its discrete representation \tilde{x}, \tilde{y} .

1890 Assuming $T \in H^1(\Omega)$, this interpolation or sampling procedure introduces first-order truncation
 1891 error: $\|T(x, y) - T(\tilde{x}, \tilde{y})\| = \mathcal{O}(\|\Delta \mathbf{x}_{r,\text{int}}\|)$. This matches classical interpolation theory over uniform
 1892 grids and reflects the interior spatial error due to limited resolution in RL's perception of the state.
 1893

1894 **Action discretization error:** The action space is discretized with resolution Δa_r . The true continuous
 1895 control a is approximated by its discretized counterpart a_r , and thus the perturbation in the heat source
 1896 term $Q(x, y, t)$ is: $\delta Q(x, y, t) = Q(a) - Q(a_r) = \mathcal{O}(\Delta a_r)$. Since the heat diffusion equation is
 1897 linear in Q , this error translates directly to a temperature deviation of the same order: $\delta T = \mathcal{O}(\Delta a_r)$.
 1898

1899 **Temporal integration error:** The PDE simulator advances the solution over a time horizon Δt_r
 1900 using $N_t = \Delta t_r / \Delta t_w$ steps of the forward Euler scheme. Each substep incurs a local truncation
 1901 error: $\tau_t^{(n)} = \mathcal{O}(\Delta t_w)$, yielding a total accumulated temporal error: $\sum_{n=1}^{N_t} \tau_t^{(n)} = \mathcal{O}(\Delta t_r)$.
 1902

1903 Combining the above components, the total one-step prediction error induced by RL interaction with
 1904 the PDE simulator in the heat diffusion system is given by:
 1905

$$\Delta_{\text{RL}} = \underbrace{C_1 \|\Delta \mathbf{x}_{r,\text{int}}\|}_{\text{State observation}} + \underbrace{C_2 \Delta a_r}_{\text{Action discretization}} + \underbrace{C_3 \Delta t_r}_{\text{Temporal integration}}.$$

1906 This decomposition forms the basis for analyzing the interaction fidelity and prediction uncertainty
 1907 of RL agents in physical systems governed by parabolic PDEs like heat diffusion.
 1908

1909
 1910 **Refined Total Error Decomposition with Time-Scale Separation.** Based on our prior analysis of
 1911 surrogate errors in both PDE and RL components, we now present a unified decomposition of all
 1912 numerical and discretization errors. Over a single RL step, the total prediction error can be expressed
 1913 as:
 1914

$$\Delta_{\text{total}} = \underbrace{C_1 \|\Delta \mathbf{x}_r\|}_{\text{RL (state space)}} + \underbrace{C_2 \Delta a_r}_{\text{RL (action discretization)}} + \underbrace{C_3 \Delta t_r}_{\text{RL (temporal propagation)}} + \underbrace{C_4 \frac{\Delta t_r}{\Delta t_w} (\|\Delta \mathbf{x}_p\| + \Delta a_w)}_{\text{PDE surrogate error, scaled by time resolution}}. \quad (\text{C.11})$$

1915
 1916 To ensure the error stays below the acceptable uncertainty level Δy , we require: $\Delta_{\text{total}} = \mathcal{O}(\Delta y)$.
 1917 This yields the following matching constraints for all discretization parameters:
 1918

$$\|\Delta \mathbf{x}_r\| \sim \|\Delta \mathbf{x}_p\| \sim \Delta a_r \sim \Delta a_w \sim \Delta t_r \sim \Delta y, \quad \text{with } \Delta t_w \ll \Delta y. \quad (\text{C.12})$$

1919
 1920 **ρ -K Analysis.** We now analyze the prediction error rate in the heat diffusion AI4S environment.
 1921 We define a prediction error event as one in which the AI4S simulator's output state does not fall
 1922 within the correct discretized cell of the RL state space. This gives rise to a relative prediction error
 1923 rate: $\rho = 1 - \frac{\Delta y}{\text{Total Prediction Error}}$, where Δy is the size of the RL spatial grid. In a two-dimensional
 1924 spatial domain (as in heat diffusion over a cooking surface), Δy corresponds to the square of the grid
 1925 length. The total prediction error contains the following components: observation noise, numerical
 1926 surrogate error and intrinsic error propagation. Over an RL planning horizon H_r , the error grows
 1927 diffusively at rate λ_1 , yielding an error amplification: $\lambda_1 \Delta y / H_r$.
 1928

1929 Substituting all components into the definition of ρ , we obtain:
 1930

$$\rho = 1 - \frac{\Delta y}{\lambda_1 \Delta y / H_r + \Delta y + C_1 \Delta y / K^2} = 1 - \frac{1}{\lambda_1 / H_r + 1 + C_1 / K^2}.$$

1931
 1932 So in the limit of high-resolution settings, i.e., when H_r and K are large, we obtain: $\rho =$
 1933 $\mathcal{O}\left(\frac{1}{H_r} + \frac{1}{K^2}\right)$. The subsequent results on optimal computational cost allocation are summarized in
 1934 Table 1 of the main text. As this involves only straightforward calculations and scaling arguments,
 1935 we omit detailed elaboration here.
 1936

1944

C.5 CART-POLE STABILIZATION

1945

1946

C.5.1 ENVIRONMENT DESCRIPTION

1947

1948

In this task, we consider the classic inverted pendulum control problem, where the objective is to apply a horizontal force to a cart in order to maintain a pendulum in an upright and stable position.

1949

1950

The state of the system is represented by a 4-dimensional vector: $\mathbf{x} = [x \ \dot{x} \ \theta \ \dot{\theta}]^\top$, where x denotes the position of the cart, \dot{x} the velocity of the cart, θ the angle of the pendulum from the vertical (positive clockwise), and $\dot{\theta}$ the angular velocity of the pendulum. The schematic of the discretization of position and angle in the inverted pendulum environment is shown in Figure C.3.c .

1951

1952

1953

1954

1955

1956

1957

1958

1959

1960

1961

The control input is a horizontal force $u(t) \in [-F_{\max}, F_{\max}]$, which we discretize for reinforcement learning purposes. The agent selects actions from this discretized control space to influence the system dynamics. The physical parameters governing the system include: M : the mass of the cart, m : the mass of the pendulum, l : the length from the pivot to the center of mass of the pendulum, g : gravitational acceleration.

Using the Lagrangian formalism, the nonlinear dynamics of the system can be derived as (Nagendra et al., 2017):

$$\ddot{\theta} = \frac{g \sin \theta - \cos \theta \left(\frac{u+ml\dot{\theta}^2 \sin \theta}{M+m} \right)}{l \left(\frac{4}{3} - \frac{m \cos^2 \theta}{M+m} \right)}$$

$$\ddot{x} = \frac{u + ml \left(\dot{\theta}^2 \sin \theta - \ddot{\theta} \cos \theta \right)}{M+m}$$

1962

1963

1964

1965

1966

1967

1968

These equations define the state transition model and control response of the inverted pendulum system. The reinforcement learning agent interacts with this environment by issuing discrete actions at fixed time intervals, aiming to stabilize the pendulum around the upright position. The reward function is typically defined to penalize deviations from the vertical orientation.

1969

1970

1971

1972

1973

C.5.2 ERROR COUPLING BETWEEN RL AND PDE SPACES IN AI4S SYSTEMS

1974

1975

1976

1977

1978

1979

Unlike the previous three environments, the control dynamics in the inverted pendulum environment are governed by an ordinary differential equation (ODE) system. But the error analysis procedure remains analogous to that used in the earlier cases. In the following, we analyze the composition of prediction error for the inverted pendulum environment.

1980

1981

1982

1983

1984

Fine-Grained Surrogate Error in Inverted Pendulum Environments. We now analyze the surrogate error incurred at each integration step in the inverted pendulum system. Since the dynamics are governed by ordinary differential equations (ODEs), the error sources are reduced to: temporal discretization, state observation through interpolation, and action quantization. We focus on the position variable $x(t)$ of the cart and quantify its one-step prediction error.

1985

1986

1987

Temporal truncation error: The ODE describing the cart’s horizontal motion is: $\ddot{x} = f(\mathbf{x}, u)$, where $u \in [-F_{\max}, F_{\max}]$ is the external force and $\mathbf{x} = [x, \dot{x}, \theta, \dot{\theta}]^\top$ is the state vector.

1988

1989

1990

1991

Using forward Euler integration for both velocity and position: $\dot{x}^{n+1} = \dot{x}^n + \Delta t_w \ddot{x}^n$, $x^{n+1} = x^n + \Delta t_w \dot{x}^n$, the local truncation error from Taylor expansion becomes: $x^{n+1} = x^n + \Delta t_w \dot{x}^n + \frac{\Delta t_w^2}{2} \ddot{x}^n + \mathcal{O}(\Delta t_w^3)$. Then the numerical derivative satisfies: $\frac{x^{n+1} - x^n}{\Delta t_w} = \dot{x}^n + \frac{1}{2} \Delta t_w \ddot{x}^n + \mathcal{O}(\Delta t_w^2)$.

1992

1993

From the expression of \ddot{x} in the inverted pendulum dynamics, we know that \ddot{x}^n is bounded. Hence, we obtain the temporal truncation error: $\tau_{\text{time}} = \mathcal{O}(\Delta t_w)$.

1994

1995

1996

1997

State discretization error : The state is often represented in a coarsely discretized state space. Suppose the cart’s true position $x \in \mathbb{R}$ is approximated by a discrete cell center \tilde{x} . Using linear interpolation or nearest-neighbor projection, the induced error is: $|x - \tilde{x}| = \mathcal{O}(\|\Delta \mathbf{x}_p\|)$, which propagates into the dynamics via the nonlinear function $f(\mathbf{x}, u)$. Assuming Lipschitz continuity of f , this implies: $\Delta_x^{\text{state}} = \mathcal{O}(\|\Delta \mathbf{x}_p\|)$.

1998 **Action-induced control error:** The control force u is discretized into a finite set, and the applied
 1999 value u_w approximates the true value with resolution Δa . The dynamics respond linearly to changes
 2000 in u , as seen from the equation for \ddot{x} , so: $\Delta_x^{\text{action}} = \mathcal{O}(\Delta a_w)$.
 2001

2002 **Final per-step surrogate error:** Summing all components, the total one-step surrogate error in
 2003 predicting the cart position is: $\Delta_x^{(k)} = \mathcal{O}(\Delta t_w) + \mathcal{O}(\|\Delta \mathbf{x}_p\|) + \mathcal{O}(\Delta a_w)$.
 2004

2005 **One-Step Prediction Error in RL-Based Inverted Pendulum Control.** In this task, we analyze
 2006 the surrogate prediction error that arises when an RL agent interacts with the dynamical system of an
 2007 inverted pendulum. The dynamics are governed by a nonlinear second-order ordinary differential
 2008 equation (ODE), and we focus our analysis on the position variable $x(t)$, which is affected by both
 2009 control input and physical evolution. The prediction error consists of three major sources: temporal
 2010 discretization, spatial resolution (observation) error, and control discretization error.
 2011

2012 **Temporal integration error:** The ODE simulator advances the solution over a time horizon Δt_r
 2013 using $N_t = \Delta t_r / \Delta t_w$ steps of the forward Euler scheme. Each substep incurs a local truncation
 2014 error: $\tau_t^{(n)} = \mathcal{O}(\Delta t_w)$, yielding a total accumulated temporal error: $\sum_{n=1}^{N_t} \tau_t^{(n)} = \mathcal{O}(\Delta t_r)$.
 2015

2016 **Observation-induced spatial error:** Suppose the RL agent does not access the true continuous state
 2017 $x(t_n)$, but instead receives a quantized observation $\tilde{x}(t_n)$, derived from uniform discretization over
 2018 spatial grid resolution Δx_r . In one-dimensional observation with linear interpolation or rounding,
 2019 the state projection satisfies: $|x(t_n) - \tilde{x}(t_n)| \leq \frac{1}{2} \Delta x_r$. This discrepancy is formally equivalent to a
 2020 spatial interpolation error and yields a spatial state uncertainty: $\tau_{\text{state}} = \mathcal{O}(\Delta x_r)$.
 2021

2022 **Action discretization error:** The control input u is issued from a discretized action set $a_r \in \mathcal{A}_r$, such
 2023 that: $u(t_n) = \bar{u}(t_n) + \delta u(t_n)$, $\delta u = \mathcal{O}(\Delta a_r)$, where $\bar{u}(t_n)$ is the true continuous control signal
 2024 and δu is the discretization mismatch. From the expression of \ddot{x} , this control deviation contributes
 2025 linearly to the evolution of position through: $\delta x^{n+1} = \Delta t_r \delta \dot{x}^{n+1} = \Delta t_r^2 \cdot \frac{\delta u}{M+m} = \mathcal{O}(\Delta a_r)$. We
 2026 treat time as a fixed constant when integrating the accumulated effect of the action, rather than
 2027 considering it as an error term.
 2028

2029 **Total error:** Combining the above contributions from temporal, spatial, and control discretization
 2030 errors, the one-step prediction error in the inverted pendulum environment satisfies:
 2031

$$\Delta_{\text{pendulum}} = \underbrace{C_1 \Delta x_r}_{\text{State resolution}} + \underbrace{C_2 \Delta a_r}_{\text{Action resolution}} + \underbrace{C_3 \Delta t_r}_{\text{Temporal integration}} .$$

2032 **Refined Total Error Decomposition with Time-Scale Separation.** Based on our previous analysis
 2033 of surrogate error in both the ODE and RL components, we now integrate all sources of discretization
 2034 and numerical error into a unified framework. Over a single RL step, the total prediction error can be
 2035 decomposed as follows:
 2036

$$\Delta_{\text{total}} = \underbrace{C_1 \|\Delta \mathbf{x}_r\|}_{\text{RL (state space)}} + \underbrace{C_2 \Delta a_r}_{\text{RL (action discretization)}} + \underbrace{C_3 \Delta t_r}_{\text{RL (temporal propagation)}} + \underbrace{C_4 \frac{\Delta t_r}{\Delta t_w} (\|\Delta \mathbf{x}_p\| + \Delta a_w)}_{\text{PDE surrogate error, scaled by time resolution}} . \quad (\text{C.13})$$

2037 To ensure that the learned RL transition kernel remains robust under observation uncertainty Δy ,
 2038 we require that: $\Delta_{\text{total}} = \mathcal{O}(\Delta y)$. This constraint implies a resolution matching condition across all
 2039 discretization dimensions. Specifically, we obtain the following asymptotic scaling relations:
 2040

$$\|\Delta \mathbf{x}_r\| \sim \|\Delta \mathbf{x}_p\| \sim \Delta a_r \sim \Delta a_w \sim \Delta t_r \sim \Delta y, \quad \text{with } \Delta t_w \ll \Delta y. \quad (\text{C.14})$$

2041 **ρ -K Analysis.** We now analyze the prediction error rate in the cart-pole environment. The total
 2042 prediction error contains the following components: observation noise, numerical surrogate error and
 2043 intrinsic error propagation. Over an RL planning horizon H_r , the error grows diffusively at rate λ_1 ,
 2044 yielding an error amplification: $\lambda_1 \Delta y / H_r$. Substituting all components into the definition of ρ , we
 2045 obtain:
 2046

$$\rho = 1 - \frac{\Delta y}{\lambda_1 \Delta y / H_r + \Delta y + C_1 \Delta y / K} = 1 - \frac{1}{\lambda_1 / H_r + 1 + C_1 / K}.$$

2052 Since there is no notion of state grids in the inverted pendulum environment and all state discretizations
 2053 are first-order, the order of K is one. So in the limit of high-resolution settings, i.e., when H_r and
 2054 K are large, we obtain: $\rho = \mathcal{O}\left(\frac{1}{H_r} + \frac{1}{K}\right)$. The subsequent results on optimal computational
 2055 cost allocation are summarized in Table 1 of the main text. As this involves only straightforward
 2056 calculations and scaling arguments, we omit detailed elaboration here.
 2057

2058 C.5.3 EMPIRICAL COMPARISON OF TABULAR RL AND DEEP Q-LEARNING 2059

2060 We designed an experiment to test our theoretical findings on different learning algorithms. Using a
 2061 data-driven inverted pendulum (CartPole), we compared the sample complexity of Tabular RL and
 2062 Q-Learning. This comparison was performed under various discrete observation resolutions.

2063 **Experimental Setup** The CartPole system consists of a cart moving along a horizontal track with a
 2064 hinged pole on top. The agent’s goal is to apply horizontal forces to maintain the pole in an upright
 2065 position. We adopt a two-stage workflow: (i) a neural network dynamics model is trained from offline
 2066 simulation data, and (ii) this learned model serves as the environment for RL training.

2067 **Reward Function** To provide smooth learning signals and mitigate oscillations, we employ a shaped
 2068 reward function:

$$2069 r_t = \begin{cases} 1 - \alpha|\theta_t| - \beta|\dot{\theta}_t| & \text{if } |\theta_t| \leq 12^\circ \text{ and } |x_t| \leq 2.4 \\ 2070 0 & \text{otherwise} \end{cases} \quad (C.15)$$

2072 where θ_t and $\dot{\theta}_t$ are the pole angle and angular velocity, respectively. The coefficients $\alpha, \beta > 0$
 2073 penalize deviation from the vertical and excessive oscillation. An episode terminates with zero reward
 2074 if the pole falls or the cart leaves the track.

2075 **Learning Methods** We train two agents under identical discretized observation resolutions:

- 2077 • **Tabular RL:** Implemented using the value iteration algorithm, with environment dynamics
 2078 derived from physics-based simulation and interpolated to the RL observation space, without
 2079 neural network approximation.
- 2080 • **Q-Learning:** Implemented using a Deep Q-Network (DQN) with discretized state and
 2081 action spaces.

2082 Performance is measured by the number of samples required to achieve a value function error below
 2083 1% and the corresponding wall-clock runtime on a single NVIDIA RTX 4090 GPU.

2085 **Results** The detailed experimental results are presented in Table C.3 and Table C.4. The results show
 2086 that Q-Learning is significantly more sample-efficient than tabular RL, which is expected due to the
 2087 generalization capability of the neural network. However, both methods show a clear sensitivity to
 2088 the resolution parameters H_r and K . Optimal performance is not achieved at the highest or lowest
 2089 resolutions but rather at an intermediate point, which is consistent with our theoretical analysis. For
 2090 tabular RL, the optimal configuration is found at $K = 1.5$ and $\log_{H_w} H_r = 1/3$. For Q-Learning,
 2091 the optimum shifts to $K = 2.0$ and $\log_{H_w} H_r = 1/2$, suggesting that DQN can better leverage a
 2092 higher-fidelity physical model. These results strongly support our central claim that even for deep
 2093 RL methods, the resolutions of the physical simulation and the RL algorithm must be carefully
 2094 coordinated to achieve optimal performance.

2095 D FUNCTIONAL ANALYSIS TOOLS FOR PDE ERROR AND LIMIT ANALYSIS 2096

2097 In this section, we provide a collection of classical results from functional analysis and PDE theory
 2098 that underpin the error analysis, convergence behavior, and design of discretization schemes in AI4S-
 2099 RL environments governed by PDEs (Evans, 1998; E, 2011). These tools help quantify how physical
 2100 solution properties interact with discretized approximations, especially under limited regularity,
 2101 complex boundaries, and dynamic state evolution.

2103 D.1 MOTIVATION: WHY DISCRETE FORMATS AND FUNCTION SPACES MATTER IN AI4S-RL

2105 In PDE-governed AI4S-RL tasks, control and prediction rely on the accurate approximation of
 2106 continuous state evolution. However, neural networks or other function approximators inherently

2106
 2107 Table C.3: Cart-Pole System: Sample Complexity $\log_{10}(N)$ for Tabular RL to Achieve Value
 2108 Function Error < 0.01 .

$\log_{H_w} H_r$	K=1.0	K=1.5	K=2.0	K=2.5
1/8	4.51(4.58/4.45)	4.50(4.53/4.46)	4.51(4.57/4.44)	4.54(4.61/4.48)
1/6	4.26(4.32/4.19)	4.25(4.31/4.18)	4.26(4.30/4.23)	4.29(4.36/4.23)
1/5	4.10(4.17/4.04)	4.09(4.13/4.06)	4.10(4.16/4.03)	4.14(4.21/4.08)
1/4	3.91(3.98/3.85)	3.90(3.94/3.87)	3.91(3.95/3.88)	3.94(4.01/3.88)
1/3	3.77(3.81/3.74)	3.65(3.71/3.58)	3.74(3.80/3.67)	3.88(3.95/3.82)
1/2	3.95(4.02/3.89)	3.83(3.87/3.80)	3.92(3.99/3.86)	4.06(4.13/3.99)
1	4.25(4.29/4.22)	4.13(4.20/4.07)	4.22(4.26/4.19)	4.36(4.42/4.29)
2	4.55(4.62/4.49)	4.43(4.47/4.40)	4.52(4.59/4.46)	4.66(4.72/4.59)

2118
 2119 Table C.4: Cart-Pole System: Sample Complexity $\log_{10}(N)$ for Q-Learning to Achieve Value
 2120 Function Error < 0.01 .

$\log_{H_w} H_r$	K=1.0	K=1.5	K=2.0	K=2.5
1/8	3.20(3.26/3.13)	2.95(3.01/2.88)	2.82(2.86/2.79)	2.90(2.97/2.84)
1/6	3.05(3.11/2.98)	2.80(2.84/2.77)	2.72(2.78/2.65)	2.80(2.86/2.73)
1/5	2.90(2.94/2.87)	2.75(2.81/2.68)	2.68(2.72/2.65)	2.75(2.82/2.69)
1/4	2.80(2.87/2.74)	2.70(2.74/2.67)	2.67(2.73/2.61)	2.72(2.76/2.69)
1/3	2.78(2.85/2.72)	2.69(2.73/2.66)	2.66(2.72/2.59)	2.71(2.78/2.65)
1/2	2.75(2.79/2.72)	2.67(2.73/2.60)	2.65(2.69/2.62)	2.67(2.74/2.61)
1	2.85(2.91/2.78)	2.78(2.82/2.75)	2.78(2.85/2.72)	2.82(2.86/2.79)
2	3.10(3.17/3.04)	2.95(2.99/2.92)	2.93(2.98/2.87)	3.00(3.07/2.94)

2132 operate over discrete representations. Therefore, the error between the true PDE solution and the
 2133 learned approximation depends not only on the architecture but also on how the discretization aligns
 2134 with the analytical properties of the solution. Tools from functional analysis allow us to assess:

- 2136 • How regular or irregular the true solution is (e.g., near boundaries or under shocks);
- 2137 • How the discretization grid or function basis must adapt to preserve convergence;
- 2138 • How stability and generalization are affected by boundary singularities and approximation
 2139 limitations.

2141 The following theorems provide essential insights into these questions and can be naturally organized
 2142 along a logical chain used in theoretical PDE analysis:

- 2144 • Start with basic control over functions via gradients (Poincaré Inequality),
- 2145 • Link smoothness to representability and generalization (Sobolev Embedding),
- 2146 • Ensure PDE well-posedness (Lax-Milgram),
- 2147 • Connect numerical approximations to best approximants (Céa’s Lemma),
- 2148 • Use compactness to show convergence under refinement (Rellich-Kondrachov),
- 2149 • Extend to temporal-spatial regularity (Aubin-Lions),
- 2150 • Analyze boundary effects (Maximum Principle, Trace Theorem),
- 2151 • Control physical evolution through energy bounds (Energy Estimates).

2155 The materials presented in this section are well-established results in computational mathematics
 2156 and are not original contributions by the authors. Our aim is to provide readers—particularly those
 2157 interested in the error analysis of PDE-based surrogates—with an introductory overview of the
 2158 foundational knowledge in this field. We emphasize concepts essential for understanding boundary
 2159 error behavior in PDEs, especially those involving the trace theorem, which is repeatedly applied
 throughout this paper.

2160 D.2 CORE THEOREMS AND THEIR ROLES
21612162 **1. Poincaré Inequality.** Let $u \in H_0^1(\Omega)$ on a bounded domain $\Omega \subset \mathbb{R}^d$. Then:

2163
$$\|u\|_{L^2(\Omega)} \leq C \|\nabla u\|_{L^2(\Omega)}.$$

2164

2165 **Interpretation:** This inequality ensures that if the gradient of a function is small, then the function
2166 itself must also be small in the L^2 sense. It captures how global behavior can be controlled through
2167 local derivatives. **Use:** In AI4S-RL systems, this inequality provides the foundation for energy-based
2168 bounds, helping ensure that learned control policies do not cause unbounded or physically unrealistic
2169 growth in system states during training or rollout (Evans, 1998).
21702171 **2. Sobolev Embedding Theorem.** For sufficiently smooth domains and appropriate s , $H^s(\Omega) \hookrightarrow$
2172 $L^p(\Omega)$ or $C^k(\Omega)$.
21732174 **Interpretation:** This theorem allows one to infer integrability or continuity from Sobolev space
2175 membership, linking smoothness to spatial regularity. **Use:** In AI4S-RL, this theorem helps justify the
2176 use of neural approximators, showing under what conditions a function learned in a weak (Sobolev)
2177 sense can be expected to generalize as a strong or continuous function (Evans, 1998).
21782179 **3. Lax–Milgram Theorem.** Given a bounded, coercive bilinear form $a(\cdot, \cdot)$ on a Hilbert space V ,
2180 the variational problem:

2181
$$a(u, v) = f(v), \quad \forall v \in V,$$

2182

2183 has a unique solution $u \in V$.
21842185 **Interpretation:** This theorem guarantees that weak formulations of PDEs are well-posed under mild
2186 assumptions. **Use:** It is critical for framing AI4S-RL problems involving constrained learning, such
2187 as solving inverse or control problems governed by PDEs. Ensures the existence and uniqueness of
2188 target states (Evans, 1998).
21892190 **4. Céa’s Lemma.** Let u solve the variational problem and $u_h \in V_h \subset V$ be the approximation.
2191 Then:

2192
$$\|u - u_h\|_V \leq \frac{C}{\alpha} \inf_{v_h \in V_h} \|u - v_h\|_V.$$

2193

2194 **Interpretation:** This lemma asserts that the best possible error of a numerical approximation is
2195 bounded proportionally by the best projection error. **Use:** In AI4S-RL, this guides the choice of rep-
2196 resentation spaces (e.g., basis functions, neural network architectures) and ensures that approximate
2197 policies or value functions can converge quasi-optimally (Hou, 2003).
21982199 **5. Rellich–Kondrachov Compactness Theorem.** If Ω is bounded and u_n is bounded in $H^1(\Omega)$,
2200 then a subsequence converges strongly in $L^2(\Omega)$.
22012202 **Interpretation:** This compactness result ensures that bounded sequences of approximate solutions
2203 have strongly convergent subsequences. **Use:** It provides a theoretical foundation for showing
2204 convergence of learned policies or solutions in AI4S-RL as the discretization or model resolution is
2205 refined (Evans, 1998).
22062207 **6. Aubin–Lions Lemma.** Let $X_0 \hookrightarrow X \hookrightarrow X_1$ be Banach spaces with compact embedding. If u_n
2208 is bounded in $L^p(0, T; X_0)$ and $\partial_t u_n$ is bounded in $L^q(0, T; X_1)$, then u_n is relatively compact in
2209 $L^p(0, T; X)$.
22102211 **Interpretation:** This lemma bridges temporal and spatial regularity, ensuring strong convergence
2212 when time derivatives and spatial norms are controlled. **Use:** Vital in PDE-driven RL settings with
2213 temporal evolution. For instance, it helps demonstrate that learned state sequences from policy
2214 iteration converge to true continuous dynamics (Lions, 1969).
22152216 **7. Maximum Principle.** Let u solve an elliptic or parabolic PDE with suitable structure. Then:

2217
$$\max_{\bar{\Omega}} u = \max_{\partial\Omega} u.$$

2218

2219 **Interpretation:** This principle ensures that solution extrema occur on boundaries, preserving
2220 physical constraints. **Use:** In AI4S-RL, this is crucial when dealing with bounded domains and
2221 enforcing reward shaping or physical limits. Helps constrain exploration near boundary zones (Majda
2222 & Bertozzi, 2002).
2223

2214 **8. Energy Estimates.** Multiply a PDE by u or $\partial_t u$, integrate over space-time, and derive:
 2215

$$2216 \quad 2217 \quad E(t) = \int_{\Omega} (|\partial_t u|^2 + |\nabla u|^2) dx \quad \text{is bounded or decaying.}$$

2218 **Interpretation:** These estimates show that system energy either remains bounded or decreases over
 2219 time. **Use:** Critical for stability analysis of RL agents in physical systems. Ensures that learned
 2220 dynamics don't lead to energy blow-up or unrealistic oscillations (E, 2011; Majda & Bertozzi, 2002).
 2221

2222 **9. Trace Theorem.** For $u \in H^1(\Omega)$, its restriction to the boundary $\partial\Omega$ satisfies:
 2223

$$2224 \quad \gamma(u) := u|_{\partial\Omega} \in H^{1/2}(\partial\Omega).$$

2225 **Interpretation:** Functions in Sobolev spaces lose half a derivative when restricted to boundaries.
 2226 **Use:** This loss of smoothness explains why control near boundaries is harder in AI4S-RL systems
 2227 and why special handling is needed in boundary-constrained reward or control definitions (Evans,
 2228 1998).
 2229

2230 D.3 LINKING THEORY TO PRACTICE

2231 These tools collectively help answer: How should one discretize state, design control inputs, and
 2232 construct reward functionals when the solution is only partially smooth, or when boundary behavior
 2233 dominates the error? Trace and Poincaré help identify when boundary refinement is necessary;
 2234 Sobolev embedding and Rellich-Kondrachov justify the use of compact approximators; Lax-Milgram
 2235 and Céa give the mathematical foundation for why value function approximators can converge.
 2236 Understanding and incorporating these results into the AI4S-RL model design allows for principled
 2237 control over generalization, robustness, and physical consistency in the learning loop.
 2238

2239 E EXPERIMENTAL DETAILS: TEPPANYAKI THERMODYNAMIC CONTROL 2240 MODEL

2241 E.1 PROBLEM FORMULATION AND PHYSICAL MODEL

2242 E.1.1 PHYSICAL SYSTEM DESCRIPTION

2243 We construct a teppanyaki temperature control simulation based on the two-dimensional unsteady
 2244 heat conduction partial differential equation. The computational domain is a uniform square metal
 2245 plate $\Omega = [0, 1] \times [0, 1] \text{ m}^2$. The system is equipped with $H_s = 2$ power-adjustable heat sources.
 2246 These sources are fixed at spatial positions $\mathbf{x}_s^{(1)} = (0.25, 0.5)$ and $\mathbf{x}_s^{(2)} = (0.75, 0.5)$. The control
 2247 variables are the heating powers $P_j(t) \in [0, 1000] \text{ W}$ for each source.
 2248

2249 Three dishes are placed on the plate. Each dish i is modeled as a square region $\mathcal{D}_i \subset \Omega$ with side
 2250 length $L_d = 0.1 \text{ m}$. Their geometric centers are located at $\mathbf{x}_d^{(1)} = (0.4, 0.5)$, $\mathbf{x}_d^{(2)} = (0.5, 0.5)$,
 2251 and $\mathbf{x}_d^{(3)} = (0.6, 0.5)$. The target temperatures are $\{70, 80, 90\}^\circ\text{C}$, and the cooking durations are
 2252 $\{50, 55, 60\}$ seconds. The control objective is to minimize the average temperature deviation of all
 2253 dishes within their cooking windows over a total duration of $T_{total} = 60$ seconds.
 2254

2255 The regional average temperature of dish i at time t is defined as:
 2256

$$2257 \quad \bar{T}_i(t) = \frac{1}{|\mathcal{D}_i|} \iint_{\mathcal{D}_i} T(\mathbf{x}, t) d\mathbf{x} \quad (\text{E.1})$$

2258 The global loss function is defined as the cumulative mean squared error:
 2259

$$2260 \quad 2261 \quad J = \int_0^{T_{total}} \sum_{i=1}^{N_d} \mathbb{I}_{[0, t_i^{cook}]}(t) \cdot (\bar{T}_i(t) - T_i^{target})^2 dt \quad (\text{E.2})$$

2262 Here $\mathbb{I}_{[0, t_i^{cook}]}(t)$ is a time indicator function. This function ensures that temperature errors are
 2263 computed only during the dish placement period.
 2264

2268 E.1.2 GOVERNING EQUATIONS
22692270 The spatiotemporal evolution of the temperature field $T(\mathbf{x}, t)$ follows the two-dimensional inhomoge-
2271 neous heat conduction equation:

2272
$$\frac{\partial T}{\partial t} = \alpha \nabla^2 T + Q(\mathbf{x}, t), \quad \mathbf{x} \in \mathbb{R}^2, t > 0 \quad (\text{E.3})$$

2273
2274

2275 where $\alpha = 5 \times 10^{-4} \text{ m}^2/\text{s}$ is the thermal diffusivity coefficient. The boundary condition adopts an
2276 infinite domain assumption: $\lim_{|\mathbf{x}| \rightarrow \infty} T(\mathbf{x}, t) = T_\infty$. The initial condition is set to the ambient
2277 temperature $T(\mathbf{x}, 0) = T_\infty = 20^\circ\text{C}$. The heat source term $Q(\mathbf{x}, t)$ is constructed by superposing H_s
2278 Gaussian-distributed heating elements:

2279
$$Q(\mathbf{x}, t) = \sum_{j=1}^{H_s} \frac{P_j(t)}{\rho c_p \cdot 2\pi\sigma^2} \exp\left(-\frac{|\mathbf{x} - \mathbf{x}_s^{(j)}|^2}{2\sigma^2}\right) \quad (\text{E.4})$$

2280
2281

2282 Here $\sigma = 0.05 \text{ m}$ is the effective radius of each heat source. The material properties are density
2283 $\rho = 7850 \text{ kg/m}^3$ and specific heat capacity $c_p = 460 \text{ J/(kg}\cdot\text{K)}$.
22842285 E.2 MULTI-RESOLUTION REINFORCEMENT LEARNING ENVIRONMENT
2286

2287 E.2.1 SPATIOTEMPORAL DISCRETIZATION SCHEME

2288 We employ a high-precision uniform Cartesian grid to partition the physical domain Ω for bottom-
2289 layer physical evolution computation. The physical grid spacing is set to $\Delta x_{phy} = 0.025 \text{ m}$. This
2290 corresponds to $N_{phy} \times N_{phy} = 40 \times 40$ grid cells. The physical time step is $\Delta t_{phy} = 1/450 \text{ s}$
2291 $\approx 0.0022 \text{ s}$. This satisfies the CFL stability condition for explicit finite difference schemes.
22922293 The coarsening degree of the RL observation space is controlled by parameter $K \geq 1$. The RL grid
2294 size is $N_{rl} = \lfloor N_{phy}/K \rfloor$ with grid spacing $\Delta x_{rl} = K \cdot \Delta x_{phy}$. Spatial coarsening is implemented
2295 through regional averaging:

2296
$$T_{rl}^{(i,j)} = \frac{1}{K^2} \sum_{p=0}^{K-1} \sum_{q=0}^{K-1} T_{phy}^{(Ki+p, Kj+q)} \quad (\text{E.5})$$

2297
2298

2299 where $T_{phy}^{(i,j)}$ represents the temperature values on the physical grid. These values are obtained by
2300 evaluating the analytical solution at the corresponding spatial coordinates $(i\Delta x_{phy}, j\Delta x_{phy})$.
23012302 The RL decision time step Δt_{rl} is determined by parameter $y = \log_{h_w}(h_r)$, where $h_w =$
2303 $T_{total}/\Delta t_{phy}$ is the temporal horizon at the physical layer and $h_r = h_w^y$ is the horizon at the
2304 RL layer. This design constructs a hierarchical simulation structure. The underlying physical envi-
2305 ronment evolves $H = \lfloor \Delta t_{rl}/\Delta t_{phy} \rfloor$ micro-steps for each decision executed by the agent. Different
2306 combinations of parameters K and y define 25 experimental configurations. The trade-off between
2307 sample efficiency and computational complexity across these configurations is the core focus of this
2308 study.
2309

2310 E.2.2 STATE-ACTION-Reward MECHANISM

2311 The system state \mathbf{s} consists of the global temperature field $T_{rl} \in \mathbb{R}^{N_{rl} \times N_{rl}}$ on the coarsened grid at the
2312 current time, along with the normalized remaining cooking time \tilde{t}_i for each food item. Parameter K
2313 defines the coarse-graining degree of observations. Temperature readings are discretized at intervals
2314 of $\Delta T_{disc} = 2^\circ\text{C} \times K$. As K increases, the agent faces dual challenges of spatial information
2315 blurring (pixelation) and temperature quantization noise. These factors significantly intensify partial
2316 observability of the environment. The action space degrades correspondingly with the state space.
2317 Power adjustment of the two heat sources is modeled as discrete actions. The adjustment granularity
2318 is directly coupled to K through the minimum power increment $\Delta P = 50 \text{ W} \times K$. This coupling
2319 causes the number of selectable power levels to decrease significantly as K increases. Coarse grids
2320 correspond to coarse control, which aligns with physical intuition.2321 To ensure consistency and physical meaning of the reward function across different time steps Δt_{rl} ,
2322 we adopt a continuous integration form based on physical micro-steps. Within the k -th RL decision

2322 step, the instantaneous reward r_k is computed as the cumulative temperature error over all physical
 2323 micro-steps in that period:

$$2325 \quad r_k = \sum_{h=1}^H \left[- \sum_{i=1}^{Num_d} \mathbb{I}_{[0, t_i^{cool}]}(t_{k,h}) \cdot \frac{(\bar{T}_i(t_{k,h}) - T_i^{target})^2}{T_{scale}^2} \right] \cdot \Delta t_{phy} \quad (E.6)$$

2328 where $H = \lfloor \Delta t_{rl} / \Delta t_{phy} \rfloor$ is the number of physical micro-steps within a single RL step, and
 2329 $T_{scale} = 50^\circ\text{C}$ is the normalization constant. Through this approach, the cumulative reward $R =$
 2330 $\sum r_k$ remains an unbiased approximation of the original continuous control objective equation E.2,
 2331 regardless of variations in RL decision frequency. This design guarantees fairness in computational
 2332 cost comparisons across different experimental groups.

2334 E.3 TRAINING ALGORITHM AND CONVERGENCE CRITERIA

2336 E.3.1 PPO ALGORITHM IMPLEMENTATION

2338 We adopt Proximal Policy Optimization (PPO) as the core training algorithm, considering the mixed
 2339 characteristics of continuous state space and discrete action space in this problem. Compared to
 2340 off-policy algorithms, PPO constrains policy update step sizes through clipped objective functions.
 2341 This mechanism effectively prevents policy oscillations in complex non-stationary thermodynamic
 2342 environments and ensures training robustness. Both the policy network (Actor) and value network
 2343 (Critic) employ multi-layer perceptron (MLP) architectures with the global state vector s as input.
 2344 The hyperparameters are set as follows: learning rate 3×10^{-4} , discount factor $\gamma = 0.99$, and
 2345 Generalized Advantage Estimation (GAE) parameter $\lambda = 0.95$. Training proceeds synchronously
 2346 across 8 parallel environments with GPU-accelerated gradient updates.

2347 E.3.2 CONVERGENCE CRITERIA

2349 To fairly compare sample efficiency across different resolution configurations (K, y) , we adopt a
 2350 relative convergence criterion. Training is considered converged and terminated when the agent
 2351 achieves the preset performance threshold in three consecutive evaluations, where each evaluation
 2352 is based on the average over 10 complete episodes. The specific threshold is set to a theoretically
 2353 derived optimal control baseline (see Section ?? for the physical environment setup). This threshold
 2354 setting ensures that computational cost comparisons across configurations are meaningful under
 2355 comparable performance levels.

2356 E.4 EXPERIMENTAL CONFIGURATION AND RESULTS

2358 The experiment covers spatial resolution ratios $K \in \{1.0, 2.0, 4.0, 6.0, 8.0\}$ and temporal resolution
 2359 parameter $y = \log_{h_w}(h_r) \in \{1/8, 1/6, 1/4, 1/3, 1/2\}$, totaling 25 configuration points. Each con-
 2360 figuration is trained with five random seeds until reaching the convergence criterion. All experiments
 2361 are executed on the same hardware platform (Intel Xeon Gold 6530, NVIDIA RTX 4090, 256 GB
 2362 RAM).

2364 E.4.1 ε - N SCALING ANALYSIS

2366 To validate the theoretical ε - N scaling predictions, we analyze complete scaling curves for all 25
 2367 resolution configurations. Here, ε denotes the suboptimality gap defined as $\varepsilon = R^* - R_{\text{current}}$, where
 2368 R^* is the optimal policy reward derived from the analytical solution. The variable N represents
 2369 training steps.

2370 Panel (c) of Figure 6 shows ε - N scaling curves for all 23 configurations that successfully reach the
 2371 convergence threshold. Different colors represent temporal resolution y values, while color intensity
 2372 indicates spatial resolution K (darker shades for smaller K , lighter for larger K). Shaded regions
 2373 illustrate the performance range across different K values within each temporal resolution group.
 2374 Vertical dashed lines mark the mean convergence step for each y , revealing clear separation between
 2375 temporal resolution groups: finer temporal resolution (smaller y) requires fewer training steps to
 converge.

2376 E.4.2 COMPUTATIONAL COST DECOMPOSITION
2377

2378 A key contribution of our experimental analysis is the rigorous decomposition of total computational
2379 cost into its constituent components. In AI4S reinforcement learning systems, computational re-
2380 sources are consumed by two fundamentally different processes: (1) physical simulation on CPU,
2381 which involves solving or evaluating PDE-based environment dynamics, and (2) neural network
2382 operations on GPU, which includes policy inference and gradient-based training updates.

2383 This decomposition is critical because CPU and GPU computational costs scale differently with
2384 resolution parameters. CPU costs are dominated by the complexity of the physical solver, which
2385 depends on grid resolution and time step constraints imposed by numerical stability. GPU costs
2386 are primarily determined by the number of RL training iterations required for convergence, which
2387 depends on the effective state-action space size and the difficulty of credit assignment under different
2388 temporal horizons.

2389 Panels (d), (e), and (f) of Figure 6 present heatmaps showing the CPU, GPU, and total computational
2390 costs respectively across all 25 configurations. Several important observations emerge from this
2391 decomposition:

2392 **CPU-GPU trade-off structure.** The CPU and GPU cost landscapes exhibit distinct patterns. CPU
2393 costs tend to increase with finer temporal resolution (smaller $\log_{h_w}(h_r)$) due to the larger number of
2394 physical simulation steps required. GPU costs show more complex dependence on both spatial and
2395 temporal parameters, reflecting the interplay between state space complexity and horizon length in
2396 the RL training dynamics.

2397 **Non-additive composition.** The total cost is not simply the sum of CPU and GPU costs in logarithmic
2398 scale. The relative contribution of each component varies significantly across the configuration space.
2399 At fine resolutions, CPU costs dominate due to expensive physical simulations. At coarse resolutions,
2400 GPU costs become relatively more significant as the RL algorithm struggles with reduced observability
2401 and must compensate through additional training iterations.

2402 **Optimal configuration identification.** The optimal configuration ($K = 6.0, \log_{h_w}(h_r) = 1/3$)
2403 achieves the minimum total cost by balancing these competing effects. This configuration provides
2404 sufficient resolution for the agent to learn effective control policies while avoiding the computational
2405 overhead of unnecessarily fine discretization.

2407 E.4.3 COMPUTATIONAL COST SUMMARY
2408

2409 For each configuration, the logical interaction count N_{total} recorded by environment destructors
2410 is combined with the sampling multiplier $M_{\text{sample}}(K, \Delta t)$ measured from post-training sampling
2411 experiments (see Section E.5). The equivalent total computational cost is reconstructed in units of
2412 $\log_{10}(\text{FLOPs}_{\text{total}})$. Detailed numerical data with confidence intervals are presented in Table E.5.

2414 Table E.5: Equivalent Total Computational Cost ($\log_{10}(\text{FLOPs}_{\text{total}})$) with Confidence Intervals
2415

$\log_{h_w}(h_r)$	$K = 1$	$K = 2$	$K = 4$	$K = 6$	$K = 8$
1/8	17.77 ± 0.29	17.74 ± 0.35	17.72 ± 0.26	17.77 ± 0.14	R.I.
1/6	17.59 ± 0.42	17.60 ± 0.29	17.65 ± 0.21	17.65 ± 0.41	R.I.
1/4	17.59 ± 0.26	17.56 ± 0.13	17.50 ± 0.17	17.54 ± 0.19	17.81 ± 0.27
1/3	17.64 ± 0.17	17.59 ± 0.22	17.41 ± 0.20	17.27 ± 0.22	17.69 ± 0.10
1/2	17.60 ± 0.15	17.70 ± 0.19	17.58 ± 0.20	17.36 ± 0.08	18.00 ± 0.08

2423 **Note:** Values represent $\log_{10}(\text{FLOPs}_{\text{total}})$ with standard deviation computed from five random seeds.
2424 Both the logical interaction count N_{total} during training and the statistical sampling multiplier M_{sample}
2425 are comprehensively considered. Configurations marked **R.I.** (Resolution Insufficient) fail to reach the
2426 performance threshold due to excessive discretization coarseness. Specifically, $K = 8.0$ corresponds
2427 to only 5×5 RL observation grids and action resolution of $\Delta P = 400$ W (with merely 3 selectable
2428 levels: 0 W, 400 W, 800 W). Under this granularity, the agent cannot achieve fine temperature control.
2429 The optimal configuration ($K = 6.0, \log_{h_w}(h_r) = 1/3$) is highlighted in bold, corresponding to
the lowest equivalent computational complexity of $10^{17.27 \pm 0.22}$ FLOPs. This result confirms the

2430 theoretical prediction that both excessively coarse and excessively fine discretization lead to efficiency
 2431 losses.

2432 To ensure the reliability of our conclusions, we conducted extensive statistical analysis across multiple
 2433 random seeds. Each configuration was trained with five independent random seeds, controlling for
 2434 variations in network initialization, environment stochasticity, and optimization trajectory. The
 2435 confidence intervals reported in Table E.5 and visualized in Figure 6 represent the standard deviation
 2436 of the total computational cost across these runs.

2437

2438 E.5 COMPUTATIONAL COST RECONSTRUCTION METHODOLOGY

2439

2440 To ensure that experimental conclusions generalize to real AI4S scenarios where numerical PDE
 2441 solvers must be used, we adopt a computational cost reconstruction approach. The core challenge
 2442 is that directly recording the CPU/GPU execution time for every state transition during training is
 2443 complex and implementation-dependent. Specifically, each invocation of the environment’s physics
 2444 computation (via destructor calls for object lifecycle management) may require vastly different
 2445 convergence iterations depending on local condition numbers of the linear systems or nonlinear
 2446 solver states at different grid cells. Recording these fine-grained execution times would require
 2447 instrumenting every grid cell’s solver, which is impractical for large-scale experiments.

2448 Instead, we adopt a two-stage approach: first, we complete the full PPO training for all 25 resolution
 2449 configurations and record only the *logical interaction count* N_{total} via destructor counters in the
 2450 environment core class. This count reflects the sample complexity of the RL algorithm itself,
 2451 independent of the underlying physics solver implementation. Second, after training converges, we
 2452 perform a post-hoc sampling study to estimate the *computational multiplier* that would be required in
 2453 real scenarios using numerical solvers.

2454 Specifically, we randomly select a representative subset of state-action pairs encountered during
 2455 converged policy execution. For each selected state, we perform Monte Carlo sampling within the
 2456 observation uncertainty range $[\mathbf{x} - \Delta y/2, \mathbf{x} + \Delta y/2]$ to measure the frequency distribution of the
 2457 next state falling into different grid cells. According to Theorem 1 in the main text, the number of
 2458 samples $M_{\text{sample}}(K, \Delta t)$ required to statistically distinguish the correct target grid with confidence
 2459 $1 - \delta$ depends on the classification margin:

$$2460 \quad M_{\text{sample}}(K, \Delta t) = \mathcal{O} \left(\frac{\log(1/\delta)}{\min_j (\Delta p^{(j)})^2} \right) \quad (E.7)$$

2462 where $\Delta p^{(j)}$ is the probability gap between the dominant grid cell and its closest competitor in
 2463 dimension j .

2464 We then estimate the per-sample computational cost $C_{\text{solver}}(K)$ for a single PDE solver call at
 2465 resolution K . For explicit finite difference methods with grid size N_{phy}/K and CFL-constrained
 2466 time step, this cost scales as $C_{\text{solver}}(K) \sim (N_{\text{phy}}/K)^d$ per time step, where d is the spatial dimension.
 2467 The final reconstructed total cost is:

$$2469 \quad \text{FLOPs}_{\text{total}}(K, \Delta t) = N_{\text{total}} \times M_{\text{sample}}(K, \Delta t) \times C_{\text{solver}}(K) \quad (E.8)$$

2470 This reconstruction methodology ensures that reported computational costs honestly reflect what
 2471 would be required in real AI4S applications using numerical solvers, while still enabling efficient
 2472 large-scale hyperparameter searches during the training phase. The sampling multiplier M_{sample}
 2473 accounts for the statistical overhead of state disambiguation under observation noise, and the solver
 2474 cost C_{solver} accounts for the grid-resolution-dependent computational burden. Similar reconstruction
 2475 can be applied to GPU costs by estimating the per-operation GPU FLOPs and memory bandwidth
 2476 utilization, which typically differ from CPU costs by implementation-specific constants but follow
 2477 the same scaling laws.

2478 The physical environment observations in our experiments are obtained from the analytical solution
 2479 of the two-dimensional heat diffusion equation equation E.3 in free space. According to linear PDE
 2480 theory, the temperature field evolution can be expressed through spatiotemporal convolution of the
 2481 Green’s function (heat kernel) with the heat source term:

$$2483 \quad T(\mathbf{x}, t) = T_{\infty} + \int_0^t \iint_{\mathbb{R}^2} G(\mathbf{x}, t; \boldsymbol{\xi}, \tau) Q(\boldsymbol{\xi}, \tau) d\boldsymbol{\xi} d\tau \quad (E.9)$$

2484 where the heat kernel $G(\mathbf{x}, t; \boldsymbol{\xi}, \tau) = \frac{1}{4\pi\alpha(t-\tau)} \exp\left(-\frac{|\mathbf{x}-\boldsymbol{\xi}|^2}{4\alpha(t-\tau)}\right)$ describes the diffusion pattern of a
2485 unit point source in infinite space.
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537