
Ada-Diffuser: Latent-Aware Adaptive Diffusion for Decision-Making

Fan Feng^{1,3}, Selena Ge¹, Minghao Fu¹, Zijian Li³, Yujia Zheng², Zeyu Tang^{2,4},
Yingyao Hu^{5†}, Biwei Huang^{1†}, Kun Zhang^{2,3†}

¹ University of California San Diego ² Carnegie Mellon University

³ MBZUAI, ⁴ Stanford University, ⁵ Johns Hopkins University

† Equal Senior Authorship

Abstract

Recent work has framed decision-making as a sequence modeling problem using generative models such as diffusion models. Although promising, these approaches often overlook latent factors that exhibit evolving dynamics, elements that are fundamental to environment transitions, reward structures, and high-level agent behavior. Explicitly modeling these hidden processes is essential for both precise dynamics modeling and effective decision-making. In this paper, we propose a unified framework that explicitly incorporates latent dynamic inference into generative decision-making from minimal yet sufficient observations. We theoretically show that under mild conditions, the latent process can be identified from small temporal blocks of observations. Building on this insight, we introduce *Ada-Diffuser*, a causal diffusion model that learns the temporal structure of observed interactions and the underlying latent dynamics simultaneously, and furthermore, leverages them for planning and control. With a proper modular design, *Ada-Diffuser* supports both planning and policy learning tasks, enabling adaptation to latent variations in dynamics, rewards, and even recovering hidden action variables from action-free demonstrations. Extensive experiments on locomotion and robotic manipulation benchmarks demonstrate the model’s effectiveness in accurate latent inference, long-horizon planning, and adaptive policy learning.

1 Introduction

Learning and planning in partially observable environments is a fundamental challenge in building intelligent agents [1]. Recent work on casting decision-making as a generative modeling problem, taking advantage of powerful models such as transformers [2–4] and diffusion models [5–7], has achieved impressive results in a wide range of tasks. However, these methods often fail to account for hidden latent variables and their temporal dynamics, factors that are prevalent in real-world settings such as robotics [8], autonomous driving [9], healthcare [10, 11], and economics [12]. Ignoring such latent processes can result in suboptimal decision-making, particularly when the observational data does not provide full coverage of the latent factors underlying the environment’s dynamics [13–16].

Early works address partial observability in reinforcement learning (RL) and imitation learning (IL) by encoding historical observations and actions into belief states or latent embeddings, which represent a distribution over the underlying latent state [1, 17–20, 14]. Policy optimization or planning is then carried out based on these inferred belief states. However, learning such representations often requires access to the historical trajectories or data from a diverse set of environments. This can be prohibitively expensive, particularly in high-dimensional state or action spaces, posing challenges for integrating these methods into modern generative decision-making models, which typically prioritize scalability. Can we *identify* the latent factors that govern environment dynamics and rewards, and

integrate them into *scalable* generative decision-making models to enable adaptive planning and policy learning, using only *minimal observations*, while *preserving theoretical guarantees*?

In this paper, we pursue this goal by addressing two fundamental questions. First, what is the *minimum* set of observations required, *in principle*, to reliably identify the latent factors that govern the environment? Second, how can *latent identification* be effectively incorporated into generative models (e.g., diffusion models) to enable adaptive planning and policy learning? To answer the first question, we theoretically show that, under mild conditions, the latent factors at the time step t can be block-wise identified using only four surrounding observable measurements (i.e., state-action trajectories) within a small temporal window. This identification result implies that a small temporal block is sufficient to infer the latent factors in observational RL trajectories in an *online* manner.

Guided by the theoretical findings, we propose `Ada-Diffuser`, a novel *causal diffusion* framework with latent identification from temporal blocks, designed to model the data generation process of RL trajectories influenced by latent factors. To reflect the autoregressive nature of sequential decision making, we introduce a *causal denoising schedule* that aligns the denoising steps with the underlying causal structure, drawing inspiration from recent advances in autoregressive diffusion models [21–24]. For *temporal-block-wise latent identification*, during training, we propose a *denoise-then-refine* procedure that iteratively alternates between denoising the observations and refining latent estimates. This enables `Ada-Diffuser` to jointly learn a structured representation of latent variables and the corresponding observational distribution. At inference time, `Ada-Diffuser` generates actions and states while estimating latent variables in an online fashion. Since states and actions are conditioned on the latent factors, we employ a *zig-zag sampling* scheme that alternates between sampling state-action pairs and updating latent variables, ensuring consistency between generated sequences and their underlying latent dynamics.

`Ada-Diffuser` provides a unified generative framework for sequential decision-making. It is applicable to both *planning* and *policy learning* tasks by conditioning on different types of observations and adapting the conditional generative process accordingly. The framework is flexible and can accommodate various forms of latent, including ones that influence dynamics, rewards, or even represent high-level latent actions. Importantly, even in environments without explicitly designed latent variables, the block-wise latent identification mechanism improves generative modeling by implicitly capturing structured temporal dependencies.

Contributions: (1) We establish sufficient conditions under which latent factors influencing environment dynamics and rewards can be identified from short temporal windows of RL trajectories, without requiring full trajectory access or multi-environment data. (2) We develop `Ada-Diffuser`, a causal diffusion model that performs block-wise latent inference to jointly model latent contexts and observable trajectories. Unlike prior latent-augmented diffusion approaches, `Ada-Diffuser` introduces a minimal-sufficient block with backward refinement for identifiable latents and uses fully autoregressive denoising with zig-zag sampling to couple inference and generation. (3) `Ada-Diffuser` can be adapted to a wide range of decision-making tasks by conditioning on different types of observation. We empirically show the improved performance on a wide range of planning and control tasks, including 8 environments under 23 different settings.

2 Background and Related Work

In this section, we provide background and related work on diffusion-based decision-making. Additional discussions are provided in Appendix E, including related work on (1) learning latent belief states in POMDPs [1, 17, 19, 25, 26], particularly in the context of transfer, meta, and nonstationary RL/IL [13, 20, 27, 14, 20], and (2) autoregressive diffusion models [22–24, 28].

Recent advances use diffusion models as planners and policies for both RL and IL. *1. Diffusion Planner:* Diffusion-based planning leverages generative models to sample future state-action trajectories from a given state, using guidance techniques [29, 30] to encourage desirable properties such as high expected rewards. Taking Denoising Diffusion Probabilistic Models (DDPM [31])-based approaches as an example, these methods learn a generative model over expert trajectories $\bar{\rho} = f(s_0; a_0); \dots; (s_T; a_T)g$ by modeling a forward-noising process: $q(\mathbf{x}^t | \mathbf{x}^{t-1}) = \mathcal{N}(\mathbf{x}^t; \bar{\rho}_{-t} \mathbf{x}^{t-1}; (1 - \beta_t)\mathbf{I})$, and a parameterized denoising model $p_\theta(\mathbf{x}^{t-1} | \mathbf{x}^t)$ to reverse the process. Here, the superscript t denotes diffusion steps, T denotes the planning horizon, \mathbf{x}^0 is a clean subsequence sampled from the expert trajectory $\bar{\rho}$, and β_t controls the variance schedule at diffusion step t . During inference, trajectories are generated by starting from Gaussian noise and iteratively

denoising through the learned reverse process. This generation can be optionally conditioned on the initial state or other guidance signals \mathbf{y} (e.g., goals, rewards): $\hat{p}_\theta(\cdot | \mathbf{s}_0; \mathbf{y})$. *II. Diffusion Policy*: In contrast to diffusion planners, Diffusion Policy methods directly parameterize the policy $p_\theta(a | s)$ using diffusion models. For example, Diffusion Policy [6] uses a diffusion model to generate multi-step actions with expressive multimodal distributions. DPPO [7] extends this idea by modeling a two-layer MDP structure, which enables fine-tuning of diffusion-based policies in RL settings. Another line of work uses diffusion models to parameterize the policy networks for only the single current step [32–35]. Ada-Diffuser can generally accommodate both diffusion planner and policies within the same framework.

3 Latent Identification in POMDP

In this section, we seek to formally model the structure of the decision-making system by answering the following questions. First, where do the latent factors reside, and how do they influence the observable variables such as states, actions, and rewards? Second, can they be identified from demonstration data alone? We model the system that extends the standard MDP to include unobservable, time-varying latent variables that affect both the transition dynamics and the reward function. This model generalizes the contextual MDP by allowing the context to evolve stochastically over time. We then formalize the data generation process under this model using structural causal models (SCMs) [36]. Finally, we present theoretical results that characterize the minimal observational requirements for identifying the latent variables.

3.1 Latent Contextual POMDP with Time-Dependent Context

We model the latent factors using a general contextual MDP framework, where the context itself evolves over time. Formally, we define a latent time-varying contextual MDP as a tuple $\mathcal{M} = (S; A; C; T; R; \gamma)$, where S is the state space, A is the action space, C is the latent context space, $T(\mathbf{s}_t | \mathbf{s}_{t-1}; \mathbf{a}_{t-1}; \mathbf{c}_t)$ is the transition distribution, $R(\mathbf{s}_t; \mathbf{a}_t; \mathbf{c}_t)$ is the reward function, and $\gamma \in [0; 1]$ is the discount factor. The latent context $\mathbf{c}_t \in C$ follows a time-dependent (possibly stochastic) process: $\mathbf{c}_t \sim p(\mathbf{c}_t | \mathbf{c}_{t-1})$, and is unobserved during training and inference. The agent only observes trajectories $\mathcal{D} = \{(\mathbf{s}_0; \mathbf{a}_0); \dots; (\mathbf{s}_T; \mathbf{a}_T)\}$, and infers the latent context \mathbf{c}_t from the observational data. This is naturally relevant to several MDP models, including (dynamic) hidden parameter MDPs [37, 38, 14], Bayes-adaptive MDPs [39, 40, 13], and factored MDPs [41]. A full comparison and analysis is given in App. C.

Given trajectories generated under this model, we can describe the data generation process using the SCMs. Without the loss of generality, we consider the setting where an expert policy π^* is assumed to generate the actions, as is standard in learning from demonstration data. The data generation process can therefore be expressed as (l.h.s. Fig. 1):

Latent Dynamics: $\mathbf{c}_t = h(\mathbf{c}_{t-1}, \epsilon_t)$,

State Transitions: $\mathbf{s}_t = f(\mathbf{s}_{t-1}, \mathbf{a}_{t-1}, \mathbf{c}_t, \epsilon_t)$,

Action Generation: $\mathbf{a}_t = \pi(\mathbf{s}_t, \mathbf{c}_t)$,

Reward Function: $r_t = g(\mathbf{s}_t, \mathbf{a}_t, \mathbf{c}_t, \epsilon_t)$,




Figure 1: (a) SCM of the Latent Contextual POMDP. Gray/white nodes are observed/latent variables; green/red edges represent transitions driven by latents/expert policies, respectively. (b) Examples where latents influence either dynamics or rewards (affecting optimal actions).

where ϵ_t , ϵ_t , and ϵ_t denote i.i.d. exogenous noise variables. Fig. 1(a) shows the graphical model. Fig. 1(b) illustrates examples where latent factors on dynamics (e.g., external wind in locomotion) and rewards (e.g., varying target objects in robot control) influence optimal decisions.

3.2 Identifiability of Latent Factors with Minimal Measurements

To learn accurate dynamics and make reliable decisions, it is essential that the underlying latent factors influencing the environment are identifiable with observational data. We present theoretical results that characterize the minimal number of consecutive observations required for the identifiability of the latent variables, under a set of mild and natural assumptions.

Assumption 1 (First-order MDP). *We consider the following conditions:*

$$P(\mathbf{s}_t; \mathbf{a}_t; r_t; \mathbf{c}_t | \mathbf{s}_{t-1}; \mathbf{a}_{t-1}; \mathbf{c}_{t-1}; \epsilon_{<t-1}) = P(\mathbf{s}_t; \mathbf{a}_t; r_t; \mathbf{c}_t | \mathbf{s}_{t-1}; \mathbf{a}_{t-1}; \mathbf{c}_{t-1});$$

where $\mathbf{!}_{<t-1} = \{\mathbf{s}_{t-2}, \dots, \mathbf{s}_1; \mathbf{a}_{t-2}, \dots, \mathbf{a}_1; \mathbf{c}_{t-2}, \dots, \mathbf{c}_1\}$.

This is naturally satisfied under our setting described in Section 3.1.

Assumption 2 (Distributional Variability). *There exist observed state and action variables \mathbf{x}_t such that for any $\mathbf{x}_t \in X_t$, there exists a corresponding $\mathbf{x}_{t-1} \in X_{t-1}$ and a neighborhood N^r around $(\mathbf{x}_t; \mathbf{x}_{t-1})$ satisfying that, for all $\mathbf{x}_{t-2} \in X_{t-2}$, $\mathbf{x}_{t-1} \in X_{t-1}$, $\mathbf{x}_t \in X_t$, and $\mathbf{x}_{t+1} \in X_{t+1}$, the following conditional distribution operators are injective: (i) $L_{\mathbf{x}_{t-2}|\mathbf{x}_{t-1}}$, (ii) $L_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}$, and (iii) $L_{\mathbf{x}_t|\mathbf{x}_{t-2}, \mathbf{x}_{t-1}}$, where the conditional operator L represents transformations at the distribution level, that is, how one probability distribution is pushed forward to another [42].*

Assumption justification. Conceptually, the injectivity of these operator L implies that different inputs induce different output distributions, thus imposing a minimal condition on distributional variability. In RL systems, this condition is naturally satisfied in most stochastic environments where transitions produce sufficient diversity across different states and actions. The assumption also aligns with the conditions in identifiability theory, particularly in works using spectral decomposition and latent variable models [43–45].

Assumption 3 (Uniqueness of Spectral Decomposition). *For any $\mathbf{x}_t \in X_t$ and any $\mathbf{c}_t \in \mathcal{C}_t \subset C_t$, there exists a $\mathbf{x}_{t-1} \in X_{t-1}$ and corresponding neighborhood N^r satisfying Assumption 2 such that, for some $(\mathbf{x}_t; \mathbf{x}_{t-1}) \in N^r$ with $\mathbf{x}_t \in X_t$, $\mathbf{x}_{t-1} \in X_{t-1}$:*

- i. $0 < k(\mathbf{x}_t; \mathbf{x}_t; \mathbf{x}_{t-1}; \mathbf{x}_{t-1}; \mathbf{c}_t) < C < 1$ for any $\mathbf{c}_t \in \mathcal{C}_t$ and some constant C ;
- ii. $k(\mathbf{x}_t; \mathbf{x}_t; \mathbf{x}_{t-1}; \mathbf{x}_{t-1}; \mathbf{c}_t) \neq k(\mathbf{x}_t; \mathbf{x}_t; \mathbf{x}_{t-1}; \mathbf{x}_{t-1}; \mathbf{e}_t)$, where

$$k(\mathbf{x}_t; \mathbf{x}_t; \mathbf{x}_{t-1}; \mathbf{x}_{t-1}; \mathbf{c}_t) = \frac{\rho_{\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{c}_t}(\mathbf{x}_t | \mathbf{x}_{t-1}; \mathbf{c}_t) \rho_{\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{c}_t}(\mathbf{x}_t | \mathbf{x}_{t-1}; \mathbf{c}_t)}{\rho_{\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{c}_t}(\mathbf{x}_t | \mathbf{x}_{t-1}; \mathbf{c}_t) \rho_{\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{e}_t}(\mathbf{x}_t | \mathbf{x}_{t-1}; \mathbf{e}_t)}. \quad (1)$$

Assumption justification. Conceptually, Assumption 3 requires that k , which captures second-order variations in transition dynamics at time $t-1$ and t under the latent variable \mathbf{c} , yields distinct values for different \mathbf{c} 's. This requirement is typically met in RL, as varied latent dynamics or rewards often cause significant, observable shifts in behavior. *Crucially, this variability is precisely what motivates the need for the identification of the latent variable \mathbf{c}_t , as it governs meaningful differences in learning underlying decision-making process.*

These assumptions are mild and natural. While Assumption 1 is standard in RL, it can be relaxed without violating our theory (App. B.3.4). Assumptions 2–3 are naturally satisfied in practice, as they simply formalize that latent variables influence the dynamic, motivating why we need the identification of them. Further validation and discussion are provided in App. A.4. Under these assumptions, we establish an identifiability theory that characterizes the conditions under which the latent factors can be recovered, and specifies the level of identifiability that can be achieved.

Theorem 1 (Identifiability on Latent Factors). *Under Assumptions 1-3, the posterior distribution of latent factor with consecutive observations $p(\mathbf{c}_t | \mathbf{x}_{t-2:t+1})$ can be identifiable up to an invertible transformation on the latents $\hat{\mathbf{c}}_t = h(\mathbf{c}_t)$, where $\hat{\mathbf{c}}_t$ is estimated latents and h is an invertible function.*

The proof is in App. B.2. Theorem 1 indicates that a **short temporal window of observations (with future frame at $t+1$)** contains *sufficient* information to recover the posterior distribution over the true latent factors (up to an invertible transformation) in an online manner, without requiring access to the full trajectory. We further discuss the implications of this finding in greater detail in App. A.4.

4 Latent-Aware Adaptive Diffusion Planner and Policy

Building on Theorem 1, we introduce the `Ada-Diffuser` framework for learning and planning with latent identification. As illustrated in Fig. 2, `Ada-Diffuser` models the trajectory generation process via two modules: (1) **latent factor identification block**, which estimates the sequence of latent variables from the observable trajectories; and (2) **causal diffusion model**, which learns the causal generative process of RL trajectories and explicitly infers latent context. Guided by the theoretical findings in Theorem 1 and the generative process (Sec C), we use autoregressive denoising for temporal dependencies and a backward-refinement step over a minimal-sufficient block, designed via a tailored noise schedule and zig-zag sampling, to recover the latent posterior in an online manner.

In this section, we first present a general formulation of conditional diffusion modeling with latent variables. We then describe the two modules of `Ada-Diffuser` in detail (Fig. 2). The complete algorithmic pseudocode of the training and inference procedures are given in App. D.1.

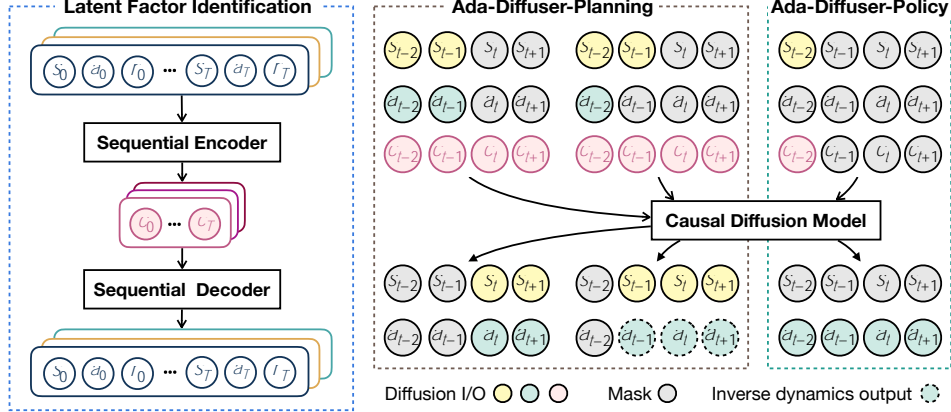


Figure 2: Overview of the Ada-Diffuser framework. The modular design consists of two main stages: latent context identification (Stage 1, Section 4.2), followed by a causal diffusion model (Stage 2, Section 4.3) that models the generative structure of the trajectories. The learned model is then used for planning or policy learning conditioned on the inferred latent context.

4.1 Latent-Augmented Diffusion Model for Planning and Policy Learning

Without loss of generality, we denote the observable trajectory as x , which may correspond to a state-action sequence s_a or a state-only sequence s , depending on the task setting. To incorporate latent structure, we augment the observable trajectory with the estimated latent context, yielding the full trajectory representation $\tilde{x} = [x; c]$, where c denotes the inferred sequence of latent variables.

We train a conditional diffusion model to generate trajectories conditioned on desired attributes $\mathcal{Y}(\cdot)$ (e.g., reward or goal specification) and the identified c . The denoising model θ is trained to predict the noise added during the forward diffusion process via the objective: $L_{\text{diff}} = \mathbb{E}_{o, y, t, k} \|\theta(\tilde{x}_t; t; \mathcal{Y}(\cdot); c) - k\|^2$, where o is a clean trajectory sample, $\tilde{x}_t \sim \mathcal{N}(0; \mathbf{I})$, and the noisy trajectory at diffusion step t is constructed as: $\tilde{x}_t = \rho_t o + \sqrt{1 - \rho_t} \tilde{x}_t$, where ρ_t denotes the cumulative product of the forward noise schedule. Here, the superscript t indexes diffusion steps, and should not be confused with the environment time step indices within the trajectory.

Ada-Diffuser can flexibly adapt to generate different components of the trajectory depending on the task. In the *planning* setting, the model generates full trajectories $\tilde{x} = \tilde{f}\mathbf{x}_t; \mathbf{x}_{t+1}; \dots; \mathbf{x}_{t+T_p}g$, where T_p denotes the planning horizon. Here, \mathbf{x}_t may have two cases: (i) $\mathbf{x}_t = \tilde{f}s_t; \mathbf{a}_t g$, when both states and actions are generated, (ii) $\mathbf{x}_t = \tilde{f}s_t g$, when only states are generated. In the latter case, we train an inverse dynamics model (IDM) [46] to infer the corresponding actions from state transitions. In the *policy learning* setting, the model generates only actions, i.e., $\tilde{x} = \tilde{f}\mathbf{a}_{t+1}; \mathbf{a}_{t+2}; \dots; \mathbf{a}_{t+T_a}g$, where T_a is the action generation horizon. While multi-step action generation methods (e.g., DP [6]) can also be viewed as a form of planning [47], for generality, we categorize such settings under the policy framework. Ada-Diffuser-Policy accommodates both variants: multi-step action generation ($T_a > 1$), as in DP, and single-step decision-making ($T_a = 1$), as in IDQL [33].

4.2 Stage 1: Offline Latent Factor Identification

Based on Theorem 1, we structure the latent inference process around *temporal blocks*, using short segments of trajectories to identify the latent context at each time step. We adopt a variational inference framework [48] in which the latent variable c_t is inferred block-wise. The prior distribution is conditioned on the latent variable from the previous step and the in-block history, while the posterior additionally incorporates future observations. Specifically, given a trajectory block $t \in T_x : t + 1$, where T_x is the block size, we have prior $p_\phi(c_t | c_{t-1})$, and posterior $q_\psi(c_t | \mathbf{x}_{t-T_x:t+1})$, where \mathbf{x} denotes the observed variables and may correspond to $\tilde{f}s g$, $\tilde{f}s; \mathbf{a} g$, or $\tilde{f}s; \mathbf{a}; r g$. We then optimize the evidence lower bound (ELBO) of the observed trajectories:

$$L_{\text{ELBO}, t} = \mathbb{E}_q(c_t | \mathbf{x}_{t-T_x:t+1}) [\log p_\theta(\mathbf{x}_t | \mathbf{x}_{t-1}; c_t)] + D_{\text{KL}}(q_\psi(c_t | \mathbf{x}_{t-T_x:t+1}) \| p_\phi(c_t | c_{t-1}))$$

Here, the reconstruction term, $\log p_\theta(\mathbf{x}_t | \mathbf{x}_{t-1}; c_t)$ is instantiated based on the available observation modalities. Specifically, (i) when only states are observed, the model reconstructs s_t conditioned

on $(\mathbf{s}_{t-1}; \mathbf{c}_t)$; and (ii) when rewards are available, the model also reconstructs r_t from $(\mathbf{s}_t; \mathbf{a}_t; \mathbf{c}_t)$. The stage is learned through a sequential encoder and decoder (l.h.s., Fig. 2).

4.3 Stage 2: Causal Diffusion Model

We propose a **causal diffusion model** for learning the generative process described in Sec.3.1. By “causal,” we refer to the modeling of the true underlying data generation process, which incorporates two key desiderata: (1) the *autoregressive process* inherent in temporal sequential RL trajectories; and (2) the *latent factor process*, capturing the causal influence of the unobserved context variables \mathbf{c}_t on the observations (e.g., $\mathbf{x}_t = [\mathbf{s}_t; \mathbf{a}_t; r_t]$). Thus, unlike prior diffusion-based RL methods and latent-augmented variants (Sec. 2; see Table A11 for a comparison), our approach incorporates the following design choices.

Autoregressive Denoising To model the autoregressive structure of trajectory generation, and following the recent advances in autoregressive diffusion [22, 23, 28], we introduce a **causal denoising schedule**. Under this mechanism, each time step within a local temporal block is assigned a denoising schedule that depends both on its temporal distance from the conditioning anchor and on the inferred latent variables. This reflects the intuition that later time steps exhibit higher uncertainty. Specifically, for a trajectory of length T , we assign monotonically increasing noise levels $\{k_1; \dots; k_T\}$, sampled linearly as $k_i = \frac{i}{T}K$ where $i \in \{1; \dots; T\}$ and K denotes the maximum diffusion step.

Given the inferred latent context $\hat{\mathbf{c}}_{0:T}$, the model performs autoregressive denoising over the block in T steps. The overall denoising process is defined as:

$$p_\theta(\mathbf{x}_0^0; \dots; \mathbf{x}_{T-1}^0 \mid \mathbf{x}_0^{k_1}; \dots; \mathbf{x}_{T-1}^{k_T}; \hat{\mathbf{c}}_{0:T}) \quad (2)$$

where $\mathbf{x}_i^{k_i}$ denotes the noisy observation at time step i , and \mathbf{x}_i^0 is the clean, denoised output.

Specifically, the first denoising step is: $p_\theta(\mathbf{x}_0^0; \mathbf{x}_1^{k_1}; \dots; \mathbf{x}_{T-1}^{k_T} \mid \mathbf{x}_0^{k_1}; \dots; \mathbf{x}_{T-1}^{k_T}; \hat{\mathbf{c}}_{0:T})$; where the first observation \mathbf{x}_0 has been fully denoised and other observations are partially denoised, followed by the second step: $p_\theta(\mathbf{x}_1^0; \mathbf{x}_2^{k_2}; \dots; \mathbf{x}_{T-1}^{k_T} \mid \mathbf{x}_0^0; \mathbf{x}_1^{k_1}; \dots; \mathbf{x}_{T-1}^{k_T}; \hat{\mathbf{c}}_{0:T})$, and finally until all observations are denoised: $p_\theta(\mathbf{x}_{T-1}^0 \mid \mathbf{x}_0^0; \dots; \mathbf{x}_{T-2}^0; \mathbf{x}_{T-1}^{k_1}; \hat{\mathbf{c}}_{0:T})$.

Denoise-and-refine Mechanism Theorem 1 indicates that both historical and future observations are required for recovering the latents. However, these future observations are not accessible during online inference, which results in a mismatch between identifiability requirements and available information. Hence, guided by this insight with preserving the causal structure of the generative process, we propose a novel *denoise-and-refine mechanism* that alternates between denoising the observable sequences and refining the latent estimates, and is applied consistently during both training and inference to ensure high-quality latent context recovery in an online manner. We introduce how we implement this during training and inference.

Training: Given a noisy input $\mathbf{x}_t^{k_t}$ with noise level k_t , we first sample an initial latent context from the prior: $\hat{\mathbf{c}}_t^{\text{prior}} \sim p_\phi(\mathbf{c}_t \mid \mathbf{c}_{t-1})$, and use it to denoise the observation: $\hat{\mathbf{x}}_t^{(0)} = p_\theta(\mathbf{x}_t^{k_t}; k_t; \hat{\mathbf{c}}_t^{\text{prior}})$. Then we infer the latent using the posterior network, conditioned on a broader temporal window including future observations (accessible in offline data): $\hat{\mathbf{c}}_t^{\text{post}} \sim q_\psi(\mathbf{c}_t \mid \mathbf{x}_{t-k:t+1})$, and obtain a **refined** denoised prediction: $\hat{\mathbf{x}}_t^{(0)'} = p_\theta(\mathbf{x}_t^{k_t}; k_t; \hat{\mathbf{c}}_t^{\text{post}})$.

We have two reconstruction losses: one from the prior-sampled latent, $L_{\text{prior}} = k\hat{\mathbf{x}}_t^{(0)} - \mathbf{x}_t^0 k^2$, and one from the posterior-sampled latent, $L_{\text{post}} = k\hat{\mathbf{x}}_t^{(0)'} - \mathbf{x}_t^0 k^2$. To encourage the posterior latent to produce better reconstructions, we introduce a **contrastive improvement loss**: $L_{\text{rel}} = \text{softplus}(\log L_{\text{post}} - \log \text{sg}(L_{\text{prior}} + m))$, where $\text{sg}(\cdot)$ denotes stop-gradient, $\text{softplus}(u) = \log(1 + e^u)$, and $m > 0$ is a margin hyperparameter. The final objective for this denoise-and-refine step is: $L_{\text{d-r}} = L_{\text{post}} + \text{prior}L_{\text{prior}} + \text{rel}L_{\text{rel}}$, where prior and rel are weighting coefficients.

Inference: During inference, future observations are not available, which prevents direct use of the posterior network for latent inference. To address this, we adopt a *zig-zag sampling strategy*¹ that

¹Note on terminology: our use of “zig-zag” is purely descriptive, and there is no connection between the proposed sampling and Bai et al. [49].




Figure 3: zig-zag sampling (2 steps).




Figure 4: (a). Identification Results (i.e., Linear Probing MSE, R^2) and normalized rewards on the Cheetah environment with time-varying wind as the latent factor, evaluated across different block sizes. (b). Results (i.e., average success rate) on planning with action-free demonstrations on Robomimic benchmark. "AF" denotes Action-free.

combines autoregressive denoising with latent refinement. Specifically, we first sample the entire trajectory by applying the forward diffusion process with the maximum noise level K . We then perform autoregressive denoising across time.

For each time step t , we begin by denoising \mathbf{x}_t^K to an intermediate noise level k_1 using $\hat{\mathbf{c}}_t$ sampled from the prior: $\mathcal{C}_t^{\text{prior}} \rho_\phi(\mathbf{c}_t | \mathbf{c}_{t-1})$. We then obtain updated $\hat{\mathbf{c}}_t$ from the posterior latent distribution $\mathcal{C}_t^{\text{post}} q_\psi(\mathbf{c}_t | \mathbf{x}_{t-k:t-1}^0; \mathbf{x}_t^{k_1}; \mathbf{x}_{t+1}^{k_2})$, which is conditioned on the denoised history, the intermediate step with noise level k_1 , and the next step with noise level k_2 . We then use $\mathcal{C}_t^{\text{post}}$ as the input to further denoise $\mathbf{x}_t^{k_1}$ to \mathbf{x}_t^0 . An illustration of the zig-zag inference process is provided in Fig. 3².

In summary, Ada-Diffuser leverages autoregressive noise scheduling to reflect temporal structure, integrates latent context identification by the *denoise-and-refine* mechanism, and employs *zig-zag sampling* for online latent inference. This framework accommodates a wide range of scenarios, including latent dynamics/rewards, learning from action-free data with latent actions, and both state- and image-based environments. All variants share the same core, with task-specific modifications to the input/output only. Details of these architectural and variations are in App. H.

5 Experiments

We aim to answer the following questions in the evaluation: (1) *Latent Identification*: How well can Ada-Diffuser capture latent factors in the environment? (2) *Learning with Latent Factors*: How effective is Ada-Diffuser in planning and control when learning with the latent context on dynamics and reward? And can Ada-Diffuser infer latent actions from action-free demonstrations? (3) *Learning with Environments w/o Explicit Latents*: In environments without explicit latent factors, can modeling latent processes still bring performance gains? (4) *Ablation Studies*: What is the impact of key design choices in the framework?

Benchmarks We consider a diverse set of benchmarks, including Mujoco-based locomotion tasks (Cheetah, Ant, Walker), a robot navigation task (Maze2D), and a robot arm control task (Franka-Kitchen) [50], all from the D4RL benchmark suite [51]. We also consider robotic manipulation tasks from RobotMimic [52] and LIBERO-10 [53]. A detailed description and illustration of these environments is provided in App. F. We introduce latent factors affecting both dynamics (\mathbf{c}^s) and reward functions (\mathbf{c}^r) in the Cheetah and Ant environments, considering two types of variations: episodic changes (E) and fine-grained, time-varying step-wise changes (S). The specific change functions for each setting are detailed in App. F.1. For evaluating latent action modeling, we follow the setup from LDP [54], using action-free, pixel-based demonstrations from the LIBERO benchmark [53]. In total, we evaluate on 8 environments with 23 settings.

Baselines We compare Ada-Diffuser with a diverse set of baselines for fair and comprehensive evaluation. (1) *Vanilla diffusion models*: For planning, we consider Diffuser [5] and DD [55]. For policy learning, we include DP and IDQL [33]. We also evaluate LDCQ [56], which learns a latent skill space and optimizes a value function conditioned on both states and latent skills. (2) *Latent context modeling*: We include MetaDiffuser [27] that learns contextual representations from multiple environments. We also consider using LILAC [14] and DynaMITE [20] which models nonstationarity in RL through latent context learning using belief states. For a fair comparison, we integrate their context modules into diffusion planners and policies as plug-in components (detailed analysis in

²A larger illustration with 4 steps are given in App. Fig. A2.




Figure 5: Results on environments without explicitly designed latent factors. Complete results are provided in App. Table A6–A9.

Environment	Diffuser	DF	DF + DynaMITE	DF + LILAC	MetaDiffuser	Ours
Cheetah-Wind-E (c^s)	-120.4 ± 12.7	-105.8 ± 9.6	-82.3 ± 8.2	-91.5 ± 7.8	-95.3 ± 7.4	-68.9 ± 7.6
Cheetah-Wind-S (c^s)	-148.5 ± 9.8	-102.0 ± 10.2	-87.2 ± 10.4	-96.7 ± 9.5	-105.6 ± 14.5	-73.5 ± 8.7
Cheetah-Vel-E (c^r)	-102.4 ± 18.2	-85.6 ± 18.3	-60.2 ± 10.8	-67.8 ± 11.0	-62.6 ± 11.1	-45.8 ± 9.5
Ant-Dir-E (c^r)	188.6 ± 39.2	195.4 ± 47.0	266.7 ± 28.1	233.6 ± 31.9	229.4 ± 32.6	285.3 ± 24.5

Table 1: Results (5 seeds) on Ada-Diffuser-Planner with latent factors that affects dynamics and rewards. c^s and c^r indicate the changes on dynamics and reward, E and S represent the episodic and time-step changes. All results are averaged over 5 random seeds.

App. H.1). (3) *Latent action modeling*: We compare with LDP [54] with action-free demonstrations for planning. In total, we compare with 9 baselines across these settings.

Architecture Choices (Details are in App. D.2) For latent factor identification, we use GRU [57] embedding with MLP layers as both prior and posterior encoders to produce Gaussian distribution over latents. For decoders, we use MLP layers. For planning and policy learning, we use UNet [58] or Transformers [59] as denoising networks and use MLPs to learn the IDM. We use VAE [48] for the visual encoders.

Results on Latent Identification To verify our identification theory, we evaluate model performance under different block sizes that contain varying amounts of temporal context. We include settings where all blocks have sufficient observations, as well as a challenging case with insufficient observations (i.e., without access to future observations). To quantify the quality of the learned latent representations, we adopt linear probing and the coefficient of determination R^2 as the evaluation metric. The results, together with normalized results, are shown in Fig. 4(a). Similarly, we also provide the clustering result in App. Fig. A6. The yellow region indicates settings with insufficient observations, resulting in lower identification results. The purple region corresponds to sufficient observations and yields relatively strong performance, and the green region reflects larger block sizes, which lead to degraded results due to redundant information or inherent difficulty for optimization. Notably, the reward is positively associated with the accuracy of latent identification, validating the importance of identifying latent factors in RL trajectories.

Results on Decision-making We consider three groups based on the kind of latent factors.

Group I: Latent factors on dynamics and reward: Table 1 presents the results of learning under latent factors that affect dynamics and rewards in locomotion tasks. To ensure a fair comparison, we implement autoregressive variants of DynaMITE and LILAC using the DF backbone. Results for the DP-backbone counterparts are provided in App. Tables A4 and are consistently worse than DF. Additional results, including using DP as backbones (Ada-Diffuser-policy), oracle variants and meta-learned versions of Ada-Diffuser that use ground-truth latents as input, are provided in App. Table A4–A5. From the results, we observe that Ada-Diffuser consistently achieves the best performance, with a significant margin over all baselines. In particular, it outperforms Diffusion planners and policies even when those models are enhanced with latent context modules such as DynaMITE and LILAC (pink area), which are most comparable to our setting. Furthermore, Ada-Diffuser outperforms DF, showing the effectiveness of our framework.

Group II: Latent Actions: Following Xie et al. [54], we consider learning from action-free demonstration data, where actions are treated as latent factors to be inferred. We adopt the same setup as

in [54], using a pre-trained visual encoder obtained via a VAE to learn the latent space from pixel observations. We then train a latent planner and an IDM using a diffusion-based approach. Unlike prior work, our diffusion-based latent planner additionally incorporates latent factors \mathbf{C} to model latent context. Importantly, we train only the planner using additional action-free demonstrations. Detailed training procedures are provided in App. G.1. Results on several tasks in Robomimic benchmark show that we can bring improvements on all tasks via modeling the latent process supplementary to the latent planner in [54]. Here, the IDM is trained solely on expert demonstrations. Complete results are provided in App. Table A3.

Group III: Environments w/o Explicitly Designed Latents: Crucially, in this scenario, the latent variable \mathbf{c} effectively serves as a form of Bayesian filtering over the observed trajectories, capturing the inherent stochasticity in the data (a more detailed discussion in App. D.3). Such variability commonly arises from system noise, expert action noise, or high-level unobserved factors. The results, shown in Fig. 5 (full results provided in App. Table A6–A9), support this interpretation. Even in environments without explicitly designed latent contexts, incorporating latent modeling allows Ada-Diffuser to achieve performance that is comparable to or better than these baselines. These findings suggest that our framework can consistently capture implicit latent process in the data, improving both trajectory modeling and downstream planning.

Ablation Studies We conduct ablation studies to evaluate the contributions of key components in our framework. For *latent factor identification*, Fig. 4(a) shows the effect of different temporal block sizes, illustrating the benefit of incorporating future observations during inference. For the *causal diffusion model*, we examine the impact of the following design choices: (i) removing the refinement step (*w/o refine*); (ii) removing zig-zag sampling (*w/o zig-zag*); (iii) replacing the causal noise schedule with a fixed noise level across time steps (vanilla diffusion) or with random noise scaling as in DF [22] (*same NS*, *random NS*). The results in Table 2 demonstrate the effectiveness of these modules in our framework in both settings: with and without explicit latent factors. Specifically, for refinement and zig-zag, we hypothesize the gains come from reducing posterior mismatch. We therefore run a latent probing test on Cheetah with changing wind and report linear-probe MSE across variants; Ada-Diffuser with both refinement and zig-zag attains the lowest error (Table A16; App. I.2.4). Additional ablations are provided in App. I.2, including full results, comparisons of alternative noise schedules beyond linear (App. I.2.2), sweeps over temporal block length (App. I.2.3), and analyses of long-horizon planning (App. I.2.5). Notably, we show that our method introduces no significant computational overhead in terms of training and sampling runtime (App. I.1).

Cases	Cheetah (\mathbf{c}^S)	LIBERO
Original	-73.5	93.4
w/o refine	-82.0	90.2
w/o zig-zag	-91.6	91.6
same NS	-89.7	85.2
random NS	-84.6	88.5

Table 2: Ablation on Design Choices on Cheetah-Wind-S (planner) and LIBERO (DP-policy).

6 Conclusions

We demonstrate that identifying latent factors from sequential observations is critical for effective decision-making. We provide theoretical results that establish conditions under which latent variables can be identified using small temporal blocks of observations. This insight enables a principled integration of latent identification into a diffusion-based generative framework, allowing us to capture the underlying causal process while maintaining scalability. Our proposed Ada-Diffuser is broadly applicable to a variety of settings, including planning and control tasks with or without explicit latent structure, and even action-free demonstrations. Results across diverse benchmarks show substantial improvements, validating the effectiveness of our method not only in environments with designed latent factors but also in general settings where latent structure is implicit but influential.

Acknowledgment

We would like to acknowledge the support from NSF Award No. 2229881, AI Institute for Societal Decision Making (AI-SDM), the National Institutes of Health (NIH) under Contract R01HL159805, and grants from Quris AI, Florin Court Capital, and MBZUAI-WIS Joint Program, and the AI Deira Causal Education project.

References

- [1] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
- [2] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.
- [3] Qinqing Zheng, Amy Zhang, and Aditya Grover. Online decision transformer. In *international conference on machine learning*, pages 27042–27059. PMLR, 2022.
- [4] Deqian Kong, Dehong Xu, Minglu Zhao, Bo Pang, Jianwen Xie, Andrew Lizarraga, Yuhao Huang, Sirui Xie, and Ying Nian Wu. Latent plan transformer for trajectory abstraction: Planning as latent space inference. *Advances in Neural Information Processing Systems*, 37:123379–123401, 2024.
- [5] Michael Janner, Yilun Du, Joshua B Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. *arXiv preprint arXiv:2205.09991*, 2022.
- [6] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, page 02783649241273668, 2023.
- [7] Allen Z. Ren, Justin Lidard, Lars Lien Ankile, Anthony Simeonov, Pulkit Agrawal, Anirudha Majumdar, Benjamin Burchfiel, Hongkai Dai, and Max Simchowitz. Diffusion policy optimization. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=mEppqHvbD2h>.
- [8] Mikko Lauri, David Hsu, and Joni Pajarinen. Partially observable markov decision processes in robotics: A survey. *IEEE Transactions on Robotics*, 39(1):21–40, 2022.
- [9] Zhiyu Huang, Chen Tang, Chen Lv, Masayoshi Tomizuka, and Wei Zhan. Learning online belief prediction for efficient pomdp planning in autonomous driving. *IEEE Robotics and Automation Letters*, 2024.
- [10] Milos Hauskrecht and Hamish Fraser. Planning treatment of ischemic heart disease with partially observable markov decision processes. *Artificial intelligence in medicine*, 18(3):221–244, 2000.
- [11] Daniel E Ehrmann, Shalmali Joshi, Sebastian D Goodfellow, Mjaye L Mazwi, and Danny Eytan. Making machine learning matter to clinicians: model actionability in medical decision-making. *NPJ Digital Medicine*, 6(1):7, 2023.
- [12] Gianluca Brero, Alon Eden, Darshan Chakrabarti, Matthias Gerstgrasser, Amy Greenwald, Vincent Li, and David C Parkes. Stackelberg pomdp: A reinforcement learning approach for economic design. *arXiv preprint arXiv:2210.03852*, 2022.
- [13] Luisa Zintgraf, Sebastian Schulze, Cong Lu, Leo Feng, Maximilian Igl, Kyriacos Shiarlis, Yarin Gal, Katja Hofmann, and Shimon Whiteson. Varibad: Variational bayes-adaptive deep rl via meta-learning. *Journal of Machine Learning Research*, 22(289):1–39, 2021.
- [14] Annie Xie, James Harrison, and Chelsea Finn. Deep reinforcement learning amidst continual structured non-stationarity. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 11393–11403. PMLR, 18–24 Jul 2021.
- [15] Gokul Swamy, Sanjiban Choudhury, J Bagnell, and Steven Z Wu. Sequence model imitation learning with unobserved contexts. *Advances in Neural Information Processing Systems*, 35:17665–17676, 2022.
- [16] Suneel Belkhale, Yuchen Cui, and Dorsa Sadigh. Data quality in imitation learning. *Advances in neural information processing systems*, 36:80375–80395, 2023.

- [17] Milos Hauskrecht. Value-function approximations for partially observable markov decision processes. *Journal of artificial intelligence research*, 13:33–94, 2000.
- [18] Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Bernardo A Pires, and Rémi Munos. Neural predictive belief representations. *arXiv preprint arXiv:1811.06407*, 2018.
- [19] Maximilian Igl, Luisa Zintgraf, Tuan Anh Le, Frank Wood, and Shimon Whiteson. Deep variational reinforcement learning for pomdps. In *International conference on machine learning*, pages 2117–2126. PMLR, 2018.
- [20] Anthony Liang, Guy Tennenholtz, Chih-wei Hsu, Yinlam Chow, Erdem Bıyık, and Craig Boutilier. Dynamite-rl: A dynamic model for improved temporal meta-reinforcement learning. *arXiv preprint arXiv:2402.15957*, 2024.
- [21] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022.
- [22] Boyuan Chen, Diego Martí Monsó, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. Diffusion forcing: Next-token prediction meets full-sequence diffusion. *Advances in Neural Information Processing Systems*, 37:24081–24125, 2024.
- [23] Desai Xie, Zhan Xu, Yicong Hong, Hao Tan, Difan Liu, Feng Liu, Arie Kaufman, and Yang Zhou. Progressive autoregressive video diffusion models. *arXiv preprint arXiv:2410.08151*, 2024.
- [24] Sand-AI. Magi-1: Autoregressive video generation at scale, 2025. URL https://static.magi.world/static/files/MAGI_1.pdf.
- [25] Karol Gregor, George Papamakarios, Frederic Besse, Lars Buesing, and Theophane Weber. Temporal difference variational auto-encoder. *arXiv preprint arXiv:1806.03107*, 2018.
- [26] Anirudh Goyal, Alex Lamb, Jordan Hoffmann, Shagun Sodhani, Sergey Levine, Yoshua Bengio, and Bernhard Schölkopf. Recurrent independent mechanisms. In *International Conference on Learning Representations (ICLR)*, 2021. URL <https://openreview.net/forum?id=mLcmd1EUxy->.
- [27] Fei Ni, Jianye Hao, Yao Mu, Yifu Yuan, Yan Zheng, Bin Wang, and Zhixuan Liang. Metadiffuser: Diffusion model as conditional planner for offline meta-rl. In *International Conference on Machine Learning*, pages 26087–26105. PMLR, 2023.
- [28] Tong Wu, Zhihao Fan, Xiao Liu, Hai-Tao Zheng, Yeyun Gong, Jian Jiao, Juntao Li, Jian Guo, Nan Duan, and Weizhu Chen. Ar-diffusion: Auto-regressive diffusion model for text generation. *Advances in Neural Information Processing Systems*, 36:39957–39974, 2023.
- [29] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- [30] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- [31] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [32] Zhendong Wang, Jonathan J Hunt, and Mingyuan Zhou. Diffusion policies as an expressive policy class for offline reinforcement learning. *arXiv preprint arXiv:2208.06193*, 2022.
- [33] Philippe Hansen-Estruch, Ilya Kostrikov, Michael Janner, Jakub Grudzien Kuba, and Sergey Levine. Idql: Implicit q-learning as an actor-critic method with diffusion policies, 2023.
- [34] Yuhui Chen, Haoran Li, and Dongbin Zhao. Boosting continuous control with consistency policy. *arXiv preprint arXiv:2310.06343*, 2023.

- [35] Cheng Lu, Huayu Chen, Jianfei Chen, Hang Su, Chongxuan Li, and Jun Zhu. Contrastive energy prediction for exact energy-guided diffusion sampling in offline reinforcement learning. In *International Conference on Machine Learning*, pages 22825–22855. PMLR, 2023.
- [36] Judea Pearl. Causal inference. *Causality: objectives and assessment*, pages 39–58, 2010.
- [37] Finale Doshi-Velez and George Konidaris. Hidden parameter markov decision processes: A semiparametric regression approach for discovering latent task parametrizations. In *IJCAI: proceedings of the conference*, volume 2016, page 1432, 2016.
- [38] Christian Perez, Felipe Petroski Such, and Theofanis Karaletsos. Generalized hidden parameter mdps: Transferable model-based rl in a handful of trials. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5403–5411, 2020.
- [39] James John Martin. *Some Bayesian decision problems in a Markov chain*. PhD thesis, Massachusetts Institute of Technology, 1965.
- [40] Michael O’Gordon Duff. *Optimal Learning: Computational procedures for Bayes-adaptive Markov decision processes*. University of Massachusetts Amherst, 2002.
- [41] Carlos Guestrin, Daphne Koller, Ronald Parr, and Shobha Venkataraman. Efficient solution algorithms for factored mdps. *Journal of Artificial Intelligence Research*, 19:399–468, 2003.
- [42] Nelson Dunford and Jacob T. Schwartz. *Linear Operators*. John Wiley & Sons, New York, 1971.
- [43] Yingyao Hu and Susanne M Schennach. Instrumental variable treatment of nonclassical measurement error models. *Econometrica*, 76(1):195–216, 2008.
- [44] Yingyao Hu and Matthew Shum. Nonparametric identification of dynamic models with unobserved state variables. *Journal of Econometrics*, 171(1):32–44, 2012.
- [45] Minghao Fu, Biwei Huang, Zijian Li, Yujia Zheng, Ignavier Ng, Guangyi Chen, Yingyao Hu, and Kun Zhang. Learning general causal structures with hidden dynamic process for climate analysis. *arXiv preprint arXiv:2501.12500*, 2025.
- [46] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua B. Tenenbaum, Tommi S. Jaakkola, and Pulkit Agrawal. Is conditional generative modeling all you need for decision making? In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=sPlfo2K9DFG>.
- [47] Zhengbang Zhu, Hanye Zhao, Haoran He, Yichao Zhong, Shenyu Zhang, Haoquan Guo, Tingting Chen, and Weinan Zhang. Diffusion models for reinforcement learning: A survey. *arXiv preprint arXiv:2311.01223*, 2023.
- [48] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *The International Conference on Learning Representations (ICLR)*, 2014.
- [49] Lichen Bai, Shitong Shao, Zikai Zhou, Zipeng Qi, Zhiqiang Xu, Haoyi Xiong, and Zeke Xie. Zigzag diffusion sampling: Diffusion models can self-improve via self-reflection. *arXiv preprint arXiv:2412.10891*, 2024.
- [50] Abhishek Gupta, Vikash Kumar, Corey Lynch, Sergey Levine, and Karol Hausman. Relay policy learning: Solving long-horizon tasks via imitation and reinforcement learning. In *Proceedings of the Conference on Robot Learning (CoRL)*. PMLR, 2020.
- [51] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.
- [52] Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. In *Conference on Robot Learning (CoRL)*, 2021.

- [53] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 36, 2023.
- [54] Amber Xie, Oleh Rybkin, Dorsa Sadigh, and Chelsea Finn. Latent diffusion planning for imitation learning. *International Conference on Machine Learning (ICML)*, 2025.
- [55] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua Tenenbaum, Tommi Jaakkola, and Pulkit Agrawal. Is conditional generative modeling all you need for decision-making? *arXiv preprint arXiv:2211.15657*, 2022.
- [56] Siddarth Venkatraman, Shivesh Khaitan, Ravi Tej Akella, John Dolan, Jeff Schneider, and Glen Berseth. Reasoning with latent diffusion in offline reinforcement learning. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=tGQirjzdd0>.
- [57] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, 2014.
- [58] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.
- [59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [60] Cuong C Nguyen, Thanh-Toan Do, and Gustavo Carneiro. Probabilistic task modelling for meta-learning. In *Uncertainty in Artificial Intelligence*, pages 781–791. PMLR, 2021.
- [61] Kate Rakelly, Aurick Zhou, Chelsea Finn, Sergey Levine, and Deirdre Quillen. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *International conference on machine learning*, pages 5331–5340. PMLR, 2019.
- [62] Annie Xie, Lisa Lee, Ted Xiao, and Chelsea Finn. Decomposing the generalization gap in imitation learning for visual robotic manipulation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3153–3160. IEEE, 2024.
- [63] Joey Hejna, Chethan Bhateja, Yichen Jiang, Karl Pertsch, and Dorsa Sadigh. Re-mix: Optimizing data mixtures for large scale imitation learning. *arXiv preprint arXiv:2408.14037*, 2024.
- [64] Jensen Gao, Annie Xie, Ted Xiao, Chelsea Finn, and Dorsa Sadigh. Efficient data collection for robotic manipulation via compositional generalization. *arXiv preprint arXiv:2403.05110*, 2024.
- [65] Nicklas Hansen, Xiaolong Wang, and Hao Su. Temporal difference learning for model predictive control. *arXiv preprint arXiv:2203.04955*, 2022.
- [66] Raymond J Carroll, Xiaohong Chen, and Yingyao Hu. Identification and estimation of nonlinear models using two samples with nonclassical measurement errors. *Journal of nonparametric statistics*, 22(4):379–399, 2010.
- [67] Assaf Hallak, Dotan Di Castro, and Shie Mannor. Contextual markov decision processes. *arXiv preprint arXiv:1502.02259*, 2015.
- [68] Shagun Sodhani, Amy Zhang, and Joelle Pineau. Multi-task reinforcement learning with context-based representations. In *International Conference on Machine Learning*, pages 9767–9779. PMLR, 2021.
- [69] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- [70] Zhe Chen et al. Bayesian filtering: From kalman filters to particle filters, and beyond. *Statistics*, 182(1):1–69, 2003.
- [71] Lawrence Rabiner and Biinghwang Juang. An introduction to hidden markov models. *iee assp magazine*, 3(1):4–16, 1986.
- [72] Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- [73] Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017.
- [74] Dean A Pomerleau. Efficient training of artificial neural networks for autonomous navigation. *Neural computation*, 3(1):88–97, 1991.
- [75] Wenhao Li. Efficient planning with latent diffusion. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=btPgDo4u4j>.
- [76] Zhixuan Liang, Yao Mu, Mingyu Ding, Fei Ni, Masayoshi Tomizuka, and Ping Luo. Adaptdiffuser: Diffusion models as adaptive self-evolving planners. *arXiv preprint arXiv:2302.01877*, 2023.
- [77] Haoran He, Chenjia Bai, Kang Xu, Zhuoran Yang, Weinan Zhang, Dong Wang, Bin Zhao, and Xuelong Li. Diffusion model is an effective planner and data synthesizer for multi-task reinforcement learning. *Advances in neural information processing systems*, 36:64896–64917, 2023.
- [78] Wei Xiao, Tsun-Hsuan Wang, Chuang Gan, Ramin Hasani, Mathias Lechner, and Daniela Rus. Safediffuser: Safe planning with diffusion probabilistic models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=ig2wk7kK9J>.
- [79] Chiyu Jiang, Andre Cornman, Cheolho Park, Benjamin Sapp, Yin Zhou, Dragomir Anguelov, et al. Motiondiffuser: Controllable multi-agent motion prediction using diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9644–9653, 2023.
- [80] Anurag Ajay, Seungwook Han, Yilun Du, Shaung Li, Abhi Gupta, Tommi Jaakkola, Josh Tenenbaum, Leslie Kaelbling, Akash Srivastava, and Pulkit Agrawal. Compositional foundation models for hierarchical planning. *arXiv preprint arXiv:2309.08587*, 2023.
- [81] Zhixuan Liang, Yao Mu, Hengbo Ma, Masayoshi Tomizuka, Mingyu Ding, and Ping Luo. Skilldiffuser: Interpretable hierarchical planning via skill abstractions in diffusion-based task execution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16467–16476, 2024.
- [82] Zibin Dong, Yifu Yuan, Jianye HAO, Fei Ni, Yao Mu, YAN ZHENG, Yujing Hu, Tangjie Lv, Changjie Fan, and Zhipeng Hu. Aligndiff: Aligning diverse human preferences via behavior-customisable diffusion model. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=bxfkIYfHyx>.
- [83] Shuang Li, Yihuai Gao, Dorsa Sadigh, and Shuran Song. Unified video action model. *arXiv preprint arXiv:2503.00200*, 2025.
- [84] Chuning Zhu, Raymond Yu, Siyuan Feng, Benjamin Burchfiel, Paarth Shah, and Abhishek Gupta. Unified world models: Coupling video and action diffusion for pretraining on large robotic datasets. *arXiv preprint arXiv:2504.02792*, 2025.
- [85] Tanmay Gangwani, Joel Lehman, Qiang Liu, and Jian Peng. Learning belief representations for imitation learning in pomdps. In *uncertainty in artificial intelligence*, pages 1061–1071. PMLR, 2020.

- [86] Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404*, 2024.
- [87] Kaifeng Gao, Jiaxin Shi, Hanwang Zhang, Chunping Wang, and Jun Xiao. Vid-gpt: Introducing gpt-style autoregressive generation in video diffusion models. *arXiv preprint arXiv:2406.10981*, 2024.
- [88] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.
- [89] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [90] Andy Shih, Suneel Belkhale, Stefano Ermon, Dorsa Sadigh, and Nima Anari. Parallel sampling of diffusion models. *Advances in Neural Information Processing Systems*, 36:4263–4276, 2023.

Appendix of Ada-Di ffuser: Latent-Aware Adaptive Diffusion for Decision-Making

3	A Discussions and Overview	2
4	A.1 Broader Impact	2
5	A.2 Limitations and Future Work	2
6	A.3 Discussions on the Core Idea	3
7	A.4 Discussions on the Theoretical Assumptions and Results	4
8	A.5 Discussions on the Model Design	5
9	A.6 Overview	6
10	B Theory	6
11	B.1 Notation List	6
12	B.2 Proof of Theorem 1	7
13	B.3 Theory-Algorithm Alignment	10
14	B.3.1 Non-Parametric Identifiability Theory	10
15	B.3.2 Model Design Guidance	10
16	B.3.3 Discussion on Assumptions	11
17	B.3.4 Relaxing Assumption 1 (beyond first-order Markov)	11
18	B.3.5 Cases in Assumption 2.	12
19	B.4 ELBO	12
20	C Summary on Different MDPs	14
21	C.1 Contextual MDPs	15
22	C.2 Hidden-Parameter MDPs	15
23	C.3 Discussions and Comparisons	15
24	D Details on Ada-Di ffuser	16
25	D.1 Full Algorithm and Results	17
26	D.2 Architecture Choices and Hyper-parameters	17
27	D.2.1 Latent Factor Identification	17
28	D.2.2 Planner	19
29	D.2.3 Policy	20
30	D.3 Connection to Bayesian Filtering	20
31	E Extended Related Works	20
32	E.1 Diffusion Model-based Decision-making	20
33	E.2 Latent Belief State Learning in POMDP	21
34	E.3 Autoregressive Diffusion Models	22

35	E.4 Summary	22
36	F Benchmark Settings and Illustrations	23
37	F.1 Latent Change Factors Design	23
38	F.2 Overview on Other Benchmarks	24
39	G Other Details on Ada-Diffuser	25
40	G.1 Latent Action Planner	25
41	G.2 Noise Scheduling	25
42	H Specific Design Choices for Baselines	25
43	H.1 Details on LILAC and DynaMITE	25
44	H.2 Details on Diffusion Forcing	26
45	I Ablation Analysis	26
46	I.1 Training/Inference Time Analysis	26
47	I.2 Ablation Results	27
48	I.2.1 Full Results Supplement to Table 2	27
49	I.2.2 Noise Schedule: Linear vs. Logistic vs. Sigmoid	27
50	I.2.3 Effect of Temporal Block Length on Latent Identification	27
51	I.2.4 Latent Probing: Effect of Backward Refinement and Zig-Zag	28
52	I.2.5 On the Effect of Planning and Execution Horizons: Long-horizon Planning	29
53	J LLM Usage Statement	34

54 **A Discussions and Overview**

55 In this section, we expand on the design and motivation behind `Ada-Diffuser`, including the
56 rationale for modeling latent factors in decision-making, key architectural choices, and additional
57 analysis of the experimental results presented in Section 5. We then provide an overview of the
58 remaining contents of this appendix.

59 **A.1 Broader Impact**

60 Our work aims to identify and leverage latent processes in generative decision-making, with applica-
61 tions in real-world domains such as robotics and healthcare. While these tasks may entail potential
62 societal risks, we do not believe any specific concerns need to be highlighted here. Instead, by uncov-
63 ering and modeling the underlying hidden processes, our approach promotes greater transparency in
64 decision-making, which can ultimately lead to more reliable and trustworthy outcomes.

65 **A.2 Limitations and Future Work**

66 One current limitation is that this work focuses primarily on theoretical formulation and algorithmic
67 development. Although we evaluate on a variety of established benchmarks, real-world deployment,

68 such as in self-driving, aerial drones, and physical robotics, remains an important direction for future
69 work.

70 A.3 Discussions on the Core Idea

Q1: On Latent Modeling. *Why is it necessary to model latent processes when we already have access to a large amount of demonstration data?*

71
72 In many decision-making systems, there exist unobservable variables that influence both the dynamics
73 and the reward structure. More generally, these latent variables often evolve over time. Such scenarios
74 are common in real-world settings, for example, in robotic control, system dynamics can be affected
75 by external forces (e.g., wind, friction), or by varying user demands (e.g., different target positions).
76 In these cases, learning an optimal policy requires conditioning on the latent factors, especially
77 when they are non-stationary or when transferring to new domains. Prior work has demonstrated the
78 importance of latent variable modeling in both reinforcement learning (RL) and imitation learning
79 (IL) [13, 20, 60, 61, 27, 14].

80 Even with access to large demonstration datasets, it remains difficult to ensure sufficient coverage
81 over the full space of environmental or task-specific latent factors relevant to decision-making. This
82 limitation has been widely acknowledged in recent efforts focused on analyzing data quality and
83 designing data collection protocols to promote generalization [16, 62–64]. However, most of these
84 works target fixed or task-specific latent variables. In contrast, we consider a more general setting
85 where latent factors evolve over time and are not predefined. Our framework provides theoretical
86 guarantees for identifying such latent variables from partial observations and seamlessly integrates
87 this identification process into diffusion models, enabling scalability across complex decision-making
88 tasks.

Q2: On the Scenarios w/o Explicit Latents. *What does latent modeling represent when no explicit latent factors are defined, and why can it still benefit decision-making?*

89
90 First, **Latent stochasticity is always present** (in real-world systems). Even in settings where all
91 task-relevant observations are available, e.g., in locomotion tasks where full physical state information
92 is provided, or in robotic manipulation with access to both proprioceptive and visual inputs, there
93 may still exist underlying processes that are not directly observed. These include domain-specific
94 factors such as external forces (e.g., wind) or dynamically changing task goals (e.g., target positions),
95 which can be viewed as implicit latent variables. Hence, it is crucial to infer and condition on these
96 latent factors

97 In the extreme case where such factors are also *fully observed*, latent modeling can still offer
98 significant benefits. Specifically, it can capture residual stochasticity present in the environment or
99 demonstration data, serving to explain variability not accounted for by observable features. As shown
100 in our formulation: $S_t = \tilde{f}(S_{t-1}; a_{t-1}; t)$; $r_t = g(S_t; a_t; t)$, the residual stochasticity ($\tilde{\cdot}$) can
101 be interpreted as implicit latent variables (sometimes can be time-correlated) influencing transitions
102 and rewards. The model can then identify meaningful structure from irrelevant or noisy variations,
103 for instance, filtering out visual background artifacts that are not predictive of dynamics or optimal
104 actions. In this sense, the learning framework is conceptually similar to Bayesian Filtering

105 Moreover, **partial observability and attribution gaps exist even in clean data**. Even in environ-
106 ments with consistent near-deterministic demonstrations, the agent often lacks access to the full
107 set of latent causal factors or attributes that influence behavior. Specifically, many systems exhibit
108 structured yet unobserved variability (e.g., task goals, preferences, intentions), and modeling this
109 variability with latent variables improves generalization.

Q3: On the Identification Theory. *What does the identification theory establish, and how does it inform algorithm design?*

110
111 The identification theory (Theorem 1) establishes that the distribution over latent variables can be
112 provably recovered from observable trajectories using only a small temporal window, specifically, a
113 small temporal block of four time steps. This provides a general non-parametric theoretical guarantee
114 that latent factors can be identified without requiring strong inductive biases or restrictive assumptions
115 on the model class or functional form.

116 This “four-step” result has direct implications for algorithm design. It suggests that latent identification
 117 can be effectively performed using a short temporal block, which aligns naturally with block-wise
 118 generative modeling approaches such as diffusion models. These models operate over segments or
 119 chunks of data, and our theoretical results justify using local temporal blocks to infer latent variables
 120 in a principled and scalable manner.

121 A.4 Discussions on the Theoretical Assumptions and Results

Q4: On the Assumptions. *What do Assumption 2 (Distributional Variability) and Assumption 3 (Uniqueness of Spectral Decomposition) mean, and why are they considered mild?*

122 We expand on the intuition and practical relevance of these two assumptions below.

123 *Distributional Variability* (Assumption 2) refers to the requirement that the conditional distributions

$$p(\mathbf{x}_{t-2} \mid \mathbf{x}_{t+1}); \quad p(\mathbf{x}_{t+1} \mid \mathbf{x}_t; \mathbf{c}_t); \quad \text{and} \quad p(\mathbf{x}_t \mid \mathbf{x}_{t-2}; \mathbf{x}_{t-1})$$

125 are sufficiently sensitive to variations in their input. That is, for different input pairs within a local
 126 neighborhood, the output distributions differ meaningfully, ensuring the system exhibits enough
 127 variability for identification. This assumption aligns with real-world decision-making settings (e.g.,
 128 locomotion or robotic manipulation), where changes in inputs such as physical state, control policy,
 129 or reward function lead to observable changes in output distributions.

130 *Uniqueness of Spectral Decomposition* (Assumption 3) builds on this by ensuring that changes in the
 131 latent variable \mathbf{c}_t induce distinct influences on the transition dynamics, specifically on the mapping
 132 from \mathbf{x}_{t-1} to \mathbf{x}_t . To formalize this, we consider the operator k :

$$k(\mathbf{x}_t; \mathbf{x}_t; \mathbf{x}_{t-1}; \mathbf{x}_{t-1}; \mathbf{c}_t) = \frac{p(\mathbf{x}_t \mid \mathbf{x}_{t-1}; \mathbf{c}_t)}{p(\mathbf{x}_t \mid \mathbf{x}_{t-1}; \mathbf{c}_t)} \frac{p(\mathbf{x}_t \mid \mathbf{x}_{t-1}; \mathbf{c}_t)}{p(\mathbf{x}_t \mid \mathbf{x}_{t-1}; \mathbf{c}_t)}, \quad (\text{A1})$$

133 which separates into two multiplicative components:

$$k_1 = \frac{p(\mathbf{x}_t \mid \mathbf{x}_{t-1}; \mathbf{c}_t)}{p(\mathbf{x}_t \mid \mathbf{x}_{t-1}; \mathbf{c}_t)}, \quad (\text{A2})$$

$$k_2 = \frac{p(\mathbf{x}_t \mid \mathbf{x}_{t-1}; \mathbf{c}_t)}{p(\mathbf{x}_t \mid \mathbf{x}_{t-1}; \mathbf{c}_t)}. \quad (\text{A3})$$

134 Here, k_1 and k_2 measure how changes in historical inputs affect the transition distribution at the
 135 current time step. The assumption requires that for any two distinct values of \mathbf{c}_t , the corresponding
 136 operator k is different, indicating that the latent variable has a sufficiently strong influence on the
 137 system dynamics.

138 Since \mathbf{x} is in the neighborhood of \mathbf{x} , this formulation effectively captures second-order changes
 139 in the transition dynamics with respect to the latent variable \mathbf{c}_t . This reflects many real-world RL
 140 systems, where even unobservable latent factors (e.g., wind speed or goal target) cause noticeable and
 141 structured changes in transition behavior over time, for instance, by considering velocity as states.

142 In summary, these two assumptions are not only theoretically necessary for identification, but also
 143 naturally hold in many RL and control systems. They justify *the need to explicitly model and identify*
 144 *latent variables*, as such variables often induce meaningful and structured changes in both dynamics
 145 and optimal decision-making behavior.

Q5: On the Identification of Posterior Distribution and up to the Invertible Function (Theorem 1). *Why do we aim to identify the posterior distribution over latent variables, and what is the role of the invertible function h between the estimated and true latents?*

146 Theorem 1 establishes that the posterior distribution over latent factors given surrounding observations,
 147 $p(\mathbf{c}_t \mid \mathbf{x}_{t-2:t+1})$, is identifiable up to an invertible transformation. That is, the estimated latent $\hat{\mathbf{c}}_t$
 148 satisfies $\hat{\mathbf{c}}_t = h(\mathbf{c}_t)$ for some invertible function h .

150 This form of identifiability is sufficient for downstream tasks such as dynamics modeling, planning,
 151 and control. Specifically, the learned dynamics or policy can be composed with h^{-1} without loss
 152 of expressiveness or utility. Since we only need to condition on the inferred latent $\hat{\mathbf{c}}_t$ to perform

153 these tasks, any invertible transformation of the latent space preserves the representational capacity
154 required for decision-making. In other words, although we may not recover the true latent variable c_t
155 exactly, the recovered representation \hat{c}_t contains the same information and can be used equivalently
156 in practice.

157 Therefore, identifying the posterior distribution (up to an invertible transformation) is both theo-
158 retically meaningful and practically sufficient for learning accurate dynamics models and optimal
159 policies.

160 A.5 Discussions on the Model Design

Q6: On Different Settings (Planning and Policy). *How is Ada-Diffuser applied to both
161 planning and policy learning settings?*

162 Ada-Diffuser is designed as a unified and generic framework that accommodates different types
163 of inputs \mathbf{x} (e.g., states, state-action pairs) and outputs (e.g., actions, trajectories, or state sequences).
164 This flexibility allows it to support a wide range of planning and policy learning paradigms. We
165 summarize four representative settings below:

- 166 • **Planning with state-action generation:** The model generates both states and actions, with
167 latent variables influencing dynamics or rewards. This setting aligns with prior work such as
168 Diffuser [5].
- 169 • **Planning with state-only generation:** The model generates future states, and an inverse
170 dynamics model is used to recover the corresponding actions. This setup follows Decision
171 Diffuser [46].
- 172 • **Planning from action-free demonstrations:** Only state sequences are available, and latent
173 variables are assumed to capture high-level behaviors or skills. This setting extends latent
174 diffusion planning [54].
- 175 • **Policy learning:** The model generates actions conditioned on the current or recent history
176 of states. This includes multi-step action generation (as in Diffusion Policy [6]) and one-step
177 action generation (as in Implicit Diffusion Q-Learning, IDQL [65]). In both cases, latent
178 factors may affect the underlying dynamics or rewards.

179 These diverse settings demonstrate the universality of our framework and highlight that uncovering
180 latent structure is a broadly applicable and critical problem in generative decision-making.

Q7: On the Latent Identification. *How is Stage 1 (Latent Identification) trained, and does it
181 introduce additional computational overhead?*

182 In Stage 1, we train the latent identification module using an offline dataset, as commonly done in
183 offline RL and imitation learning tasks. Specifically, we employ a lightweight variational autoencoder
184 (VAE) to optimize the ELBO defined in Section 4.2. Empirically, this stage introduces minimal
185 computational overhead (Appendix I.1). We further provide an ablation study in Appendix I.1
186 showing the impact of the number of training samples on the effectiveness of the latent identification
187 module.

Q8: On the Temporal Block Design. *How does this reflect Theorem 1, and why do we not use
188 exactly four steps in practice?*

189 Our approach reflects the theoretical result in Theorem 1 by identifying latent variables using small
190 temporal blocks in both Stage 1 and Stage 2. In Stage 1, we segment trajectories into local blocks
191 and optimize the ELBO to learn the posterior over latent variables. In Stage 2, we apply block-wise
192 refinement to improve the posterior estimates using both past and one-step future observations,
193 making a more accurate identification than using the prior alone.

194 While Theorem 1 shows that four consecutive time steps are sufficient for identifiability in principle,
195 we do not strictly limit the block size to four in practice. Empirically, we find that using slightly larger
196 blocks (typically between 6 and 20 steps) leads to more stable optimization and better performance.

197 Our ablations in Appendix I.2 show that without access to future observations, identifiability degrades,
 198 aligning with the theory.

199 We treat the "four-step" condition not as a strict architectural constraint but as a theoretical justification
 200 (sufficient condition) for using small temporal blocks. The optimal number of steps in practice may
 201 vary depending on data properties, task complexity, and model capacity.

Q9: On the Refinement Step. *Why is the refinement step necessary, how does it work, and does it introduce additional computational overhead?*

202
 203 The refinement step is motivated by the identification theory, which suggests that incorporating the
 204 current and future observations (other than only using historical ones) allows the model to infer
 205 a more informative posterior over latent variables than relying on the prior alone. This posterior
 206 refinement helps the model better capture latent dynamics by leveraging richer temporal context.

207 During training, the refinement step encourages the model to extract meaningful information from the
 208 posterior. Since Stage 1 optimizes the ELBO, the learned prior is already aligned with the posterior
 209 to some extent. This prevents the prior from collapsing into a trivial solution. The refinement step
 210 builds on this by using the pre-trained prior while further improving inference through contrastive
 211 learning between prior and posterior samples.

212 Importantly, this procedure does not introduce significant computational overhead. As shown in
 213 Appendix I.2, the refinement uses the same denoising network with different latent inputs (c) and
 214 adds only a lightweight contrastive loss, making it efficient in practice.

215 A.6 Overview

216 In this appendix, we first present the theoretical analysis in Section B, including the proof of Theorem 1
 217 and accompanying discussion, followed by the ELBO derivation for Ada-Diffuser. In Section C,
 218 we provide an in-depth analysis of different types of MDPs and their interconnections. Section H
 219 details the full Ada-Diffuser algorithm, model architectures, and its relation to Bayesian filtering.
 220 Section E expands on related work, covering diffusion-based decision-making, latent state estimation
 221 via belief learning, and autoregressive diffusion models. Finally, Sections F, G, H, and I provide
 222 additional details on benchmarks, baseline implementations, and complete experimental results.

223 B Theory

224 B.1 Notation List

225 We summarize the key notations used throughout the paper in Table A1, including variables for
 226 observed and latent states, temporal indices, and relevant mappings. These notations are used
 227 consistently in our theoretical analysis and algorithmic framework.

228 Also, we formally define the operators used in the following.

229 **Definition 1 (Linear Operator [42]).** *Let \mathbf{a} and \mathbf{b} be random variables with supports A and B ,
 230 respectively. The linear operator $L_{\mathbf{b}|\mathbf{a}}$ is defined as a mapping from a probability function $\rho_{\mathbf{a}} \in F(A)$
 231 to a probability function $\rho_{\mathbf{b}} \in F(B)$, given by*

$$\mathcal{F}(A) \rightarrow \mathcal{F}(B) : \rho_{\mathbf{b}} = L_{\mathbf{b}|\mathbf{a}} \circ \rho_{\mathbf{a}} = \int_{\mathcal{A}} \rho_{\mathbf{b}|\mathbf{a}}(\cdot | \mathbf{a}) \rho_{\mathbf{a}}(\mathbf{a}) d\mathbf{a} \quad (\text{A4})$$

232 Intuitively, this operator characterizes the transformation of probability distributions induced by the
 233 conditional distribution $\rho_{\mathbf{b}|\mathbf{a}}$. It provides a general representation of distributional change from \mathbf{a} to
 234 \mathbf{b} , without imposing any parametric assumptions on the underlying distributions.

235 **Definition 2 (Diagonal Operator).** *Let \mathbf{a} and \mathbf{b} be random variables with associated density
 236 functions $\rho_{\mathbf{a}}$ and $\rho_{\mathbf{b}}$ defined on supports A and B , respectively. For a fixed value $\mathbf{b} \in B$, the diagonal
 237 operator $D_{\mathbf{b}|\mathbf{a}}$ is defined as a linear operator that maps a density function $\rho_{\mathbf{a}} \in F(A)$ to a function
 238 in $F(A)$ via pointwise multiplication:*

$$D_{\mathbf{b}|\mathbf{a}} \rho_{\mathbf{a}} = \rho_{\mathbf{b}|\mathbf{a}}(\mathbf{b} | \cdot) \rho_{\mathbf{a}} \quad (\text{A5})$$

239 where $D_{\mathbf{b}|\mathbf{a}} = \rho_{\mathbf{b}|\mathbf{a}}(\mathbf{b} | \cdot)$ acts as a multiplication operator indexed by \mathbf{b} .

Index	Explanation	Support
\mathbf{x}_t	$[\mathbf{s}_t; \mathbf{a}_t]$, observed trajectories including state and action at time step t	$X_t \in \mathbb{R}^{d_a+d_s}$
d_x	dimension of observed variables	$d_a + d_s$
\mathbf{s}_t	state variable at time t	$\mathbf{s}_t \in \mathcal{S}_t$
\mathbf{a}_t	action variable at time t	$\mathbf{a}_t \in \mathcal{A}_t$
r_t	reward received at time t	$r_t \in \mathbb{R}$
\mathbf{c}_t	latent context variable at time t	$\mathbf{c}_t \in \mathcal{C}_t$
x	trajectory sequence of $(\mathbf{s}_t; \mathbf{a}_t)$	$f(\mathbf{s}_0; \mathbf{a}_0); \dots; (\mathbf{s}_T; \mathbf{a}_T)g$
c	observable trajectory (states or state-actions)	$f \text{ or } s$
	sequence of latent contexts	$\tilde{\mathbf{c}}_0; \dots; \mathbf{c}_T g$
	augmented trajectory with context	$[x; c]$
Function		
T	transition dynamics conditioned on \mathbf{c}_t	$T(\mathbf{s}_t \mathbf{s}_{t-1}; \mathbf{a}_{t-1}; \mathbf{c}_t)$
R	reward function conditioned on state, action, and context	$R(\mathbf{s}_t; \mathbf{a}_t; \mathbf{c}_t)$
E	expert policy used for generating demonstrations	$E(\mathbf{s}_t; \mathbf{c}_t)$
q_ψ	variational posterior for latent inference	$q_\psi(\mathbf{c}_t \mathbf{x}_{t-T_x:t+1})$
p_ϕ	latent prior distribution	$p_\phi(\mathbf{c}_t \mathbf{c}_{t-1})$
p_θ	generative model for transitions	$p_\theta(\mathbf{x}_t \mathbf{x}_{t-1}; \mathbf{c}_t)$
θ	denoising network in diffusion process	$\theta(\cdot)$
Symbol		
$t; t'$	exogenous noise in latent dynamics, state transitions, and reward	i.i.d. samples from noise distributions
$L_{a b}$	distribution operator from b to a	defined in Dunford and Schwartz [42]
$k(\cdot)$	ratio of joint probabilities used in uniqueness assumption	defined in Eq. 1
t	cumulative noise schedule in diffusion	product of forward noise factors
K	maximum number of diffusion steps	$K \in \mathbb{N}$
$T_p; T_a$	planning and action generation horizons	$T_p; T_a \in \mathbb{N}$
T_x	temporal block size for latent inference	$T_x \in \mathbb{N}$

Table A1: List of notations, explanations, and corresponding definitions.

240 B.2 Proof of Theorem 1

241 *Proof.* By the definition of data generation process (Fig. 1), the observed density is represented by:

$$\begin{aligned}
& p_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} \\
&= \int_{\mathcal{C}_t} \int_{\mathcal{C}_{t-1}} p_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{c}_t, \mathbf{c}_{t-1}, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} d\mathbf{c}_t d\mathbf{c}_{t-1} \\
&= \int_{\mathcal{C}_t} \int_{\mathcal{C}_{t-1}} p_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}, \mathbf{c}_t, \mathbf{c}_{t-1}} p_{\mathbf{x}_t, \mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{x}_{t-2}, \mathbf{c}_{t-1}} p_{\mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} d\mathbf{c}_t d\mathbf{c}_{t-1} \\
&= \int_{\mathcal{C}_t} \int_{\mathcal{C}_{t-1}} p_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} p_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} p_{\mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} d\mathbf{c}_t d\mathbf{c}_{t-1} \\
&= \int_{\mathcal{C}_t} \int_{\mathcal{C}_{t-1}} p_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} p_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t, \mathbf{c}_{t-1}} p_{\mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{x}_{t-2}, \mathbf{c}_{t-1}} p_{\mathbf{x}_{t-1}, \mathbf{x}_{t-2}, \mathbf{c}_{t-1}} d\mathbf{c}_t d\mathbf{c}_{t-1} \\
&= \int_{\mathcal{C}_t} \int_{\mathcal{C}_{t-1}} p_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} p_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t, \mathbf{c}_{t-1}} p_{\mathbf{c}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}, \mathbf{c}_{t-1}} d\mathbf{c}_t d\mathbf{c}_{t-1}.
\end{aligned}$$

242 Then, the property of Markov process presents conditional independence, organized as follows:

$$\begin{aligned}
p_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} &= \int_{\mathcal{C}_t} \int_{\mathcal{C}_{t-1}} p_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} p_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} p_{\mathbf{c}_t, \mathbf{c}_{t-1}, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} d\mathbf{c}_{t-1} d\mathbf{c}_t \\
&= \int_{\mathcal{C}_t} p_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} p_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} p_{\mathbf{c}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} d\mathbf{c}_t. \tag{A6}
\end{aligned}$$

243 Eq. A6 can be denoted in terms of operators: given values of $(\mathbf{x}_t; \mathbf{x}_{t-1}) \in X_t \times X_{t-1}$, Eq. A6 is

$$L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} = L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} L_{\mathbf{c}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}. \tag{A7}$$

244 Notably, Eq. A7 is the operator representation of the observed density function in 4 measurements.

245 Furthermore, the structure of Markov process implies the following two equalities:

$$\begin{aligned} \int \rho_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} &= \int_{\mathcal{C}_t} \rho_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} \rho_{\mathbf{x}_t, \mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} d\mathbf{c}_t; \\ \rho_{\mathbf{x}_t, \mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} &= \int_{\mathcal{C}_{t-1}} \rho_{\mathbf{x}_t, \mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{c}_{t-1}} \rho_{\mathbf{c}_{t-1}, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} d\mathbf{c}_{t-1}. \end{aligned} \quad (\text{A8})$$

246 For any fixed $(\mathbf{x}_t; \mathbf{x}_{t-1}) \in \mathcal{X}_t \times \mathcal{X}_{t-1}$, we notate Eq. A8 in terms of operators as follows:

$$\begin{aligned} L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} &= L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} L_{\mathbf{x}_t, \mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}; \\ L_{\mathbf{x}_t, \mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} &= L_{\mathbf{x}_t, \mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{c}_{t-1}} L_{\mathbf{c}_{t-1}, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}. \end{aligned} \quad (\text{A9})$$

247 Substituting the second line in Eq. A9 into R.H.S. of the first equation, we obtain

$$\begin{aligned} L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} &= L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} L_{\mathbf{x}_t, \mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{c}_{t-1}} L_{\mathbf{c}_{t-1}, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} \\ \int L_{\mathbf{x}_t, \mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{c}_{t-1}} L_{\mathbf{c}_{t-1}, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} &= L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}^{-1} L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}. \end{aligned} \quad (\text{A10})$$

248 The second line above uses Assumption 2 that $L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}^{-1}$ is injective. Next, we show how to
249 eliminate $L_{\mathbf{c}_{t-1}, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}$ from the above. Consider 3 measurements $(\mathbf{x}_t; \mathbf{x}_{t-1}; \mathbf{x}_{t-2})$, we have

$$\int \rho_{\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} = \int_{\mathcal{C}_{t-1}} \rho_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_{t-1}} \rho_{\mathbf{c}_{t-1}, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} d\mathbf{c}_{t-1}; \quad (\text{A11})$$

250 which, in operator notation (for fixed \mathbf{x}_{t-1}), is denoted as

$$\begin{aligned} L_{\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} &= L_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_{t-1}} L_{\mathbf{c}_{t-1}, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}; \\ L_{\mathbf{c}_{t-1}, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} &= L_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_{t-1}}^{-1} L_{\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}. \end{aligned} \quad (\text{A12})$$

251 The R.H.S. applies Assumption 2. Hence, substituting the above into Eq. A10, we obtain:

$$\begin{aligned} L_{\mathbf{x}_t, \mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{c}_{t-1}} L_{\mathbf{x}_t, \mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{c}_{t-1}}^{-1} L_{\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} &= L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}^{-1} L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} \\ L_{\mathbf{x}_t, \mathbf{c}_t | \mathbf{x}_{t-1}, \mathbf{c}_{t-1}} &= L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}^{-1} L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} L_{\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}. \end{aligned} \quad (\text{A13})$$

252 The second line applies Assumption 2 to post-multiply by $L_{\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}^{-1}$, while in the third line, we
253 postmultiply both sides by $L_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_{t-1}}$.

254 For all \mathbf{x}_t , choose a \mathbf{x}_{t-1} and a neighborhood N^r around $(\mathbf{x}_t; \mathbf{x}_{t-1})$ to satisfy Assumption 2, and pick
255 a $(\mathbf{x}_t; \mathbf{x}_{t-1})$ within the neighborhood N^r . Because $(\mathbf{x}_t; \mathbf{x}_{t-1}) \in N^r$, we also know that $(\mathbf{x}_t; \mathbf{x}_{t-1}),$
256 $(\mathbf{x}_t; \mathbf{x}_{t-1}) \in N^r$. The joint distribution of observations can be represented by Eq. A7:

$$L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} = L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} L_{\mathbf{c}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}; \quad (\text{A14})$$

257 The first term on the R.H.S., $L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}$, does not depend on \mathbf{x}_{t-1} , and the last term $L_{\mathbf{c}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}$
258 does not depend on \mathbf{x}_t . This feature suggests that, by evaluating Eq. A7 at the four pairs of points
259 $(\mathbf{x}_t; \mathbf{x}_{t-1}), (\mathbf{x}_t; \mathbf{x}_{t-1}), (\mathbf{x}_t; \mathbf{x}_{t-1}), (\mathbf{x}_t; \mathbf{x}_{t-1})$, each pair of equations will share the same operator
260 representation in common. Specifically:

$$L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} = L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} L_{\mathbf{c}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}; \quad (\text{A15})$$

$$L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} = L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} L_{\mathbf{c}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}; \quad (\text{A16})$$

$$L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} = L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} L_{\mathbf{c}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}; \quad (\text{A17})$$

$$L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} = L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} L_{\mathbf{c}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}; \quad (\text{A18})$$

261 Assumption 2 implies that $L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}$ is injective. Moreover, Assumption 3 implies $\rho_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t}(\mathbf{x}_t |$
262 $\mathbf{x}_{t-1}; \mathbf{c}_t) > 0$ for all \mathbf{c}_t , so that $D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t}$ is invertible. We can then solve for $L_{\mathbf{c}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}$ from
263 Eq. A16 as

$$D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t}^{-1} L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}^{-1} L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} = L_{\mathbf{c}_t, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}; \quad (\text{A19})$$

264 Plugging this expression into Eq. A15 leads to

$$L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} = L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} D_{\bar{\mathbf{x}}_t | \mathbf{x}_{t-1}, \mathbf{c}_t}^{-1} L_{\mathbf{x}_{t+1} | \bar{\mathbf{x}}_t, \mathbf{c}_t}^{-1} L_{\mathbf{x}_{t+1}, \bar{\mathbf{x}}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}. \quad (\text{A20})$$

265 At this point, we have decomposed the observable joint operator and expressed it in terms of latent-
266 conditioned transitions, enabling spectral analysis for identifying latent structure.

267 Lemma 1 of [43] shows that, given the injectivity of $L_{\mathbf{x}_t, \bar{\mathbf{x}}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}$ as in Assumption 2, we can
268 postmultiply by $L_{\bar{\mathbf{x}}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}^{-1}$ to obtain:

$$\mathbf{M} \quad L_{\mathbf{x}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} L_{\bar{\mathbf{x}}_{t+1}, \mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}^{-1} = L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} D_{\bar{\mathbf{x}}_t | \mathbf{x}_{t-1}, \mathbf{c}_t}^{-1} L_{\bar{\mathbf{x}}_{t+1} | \bar{\mathbf{x}}_t, \mathbf{c}_t}^{-1}. \quad (\text{A21})$$

269 Similarly, manipulations of Eq. A17 and A18 lead to

$$\mathbf{N} \quad L_{\mathbf{x}_{t+1}, \bar{\mathbf{x}}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}} L_{\bar{\mathbf{x}}_{t+1}, \mathbf{x}_t, \bar{\mathbf{x}}_t, \mathbf{x}_{t-1}, \mathbf{x}_{t-2}}^{-1} = L_{\mathbf{x}_{t+1} | \bar{\mathbf{x}}_t, \mathbf{c}_t} D_{\bar{\mathbf{x}}_t | \bar{\mathbf{x}}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \bar{\mathbf{x}}_t, \mathbf{c}_t}^{-1} L_{\bar{\mathbf{x}}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}^{-1}. \quad (\text{A22})$$

270 Assumption 2 guarantees that, for any \mathbf{x}_t , $(\mathbf{x}_t; \mathbf{x}_{t-1}; \mathbf{x}_{t-2})$ exist so that Eq. A21 and Eq. A22 are
271 valid operations. Finally, we postmultiply Eq. A21 by Eq. A22 to obtain:

$$\begin{aligned} \mathbf{MN} &= L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} D_{\bar{\mathbf{x}}_t | \mathbf{x}_{t-1}, \mathbf{c}_t}^{-1} L_{\mathbf{x}_{t+1} | \bar{\mathbf{x}}_t, \mathbf{c}_t} L_{\bar{\mathbf{x}}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}^{-1} D_{\bar{\mathbf{x}}_t | \bar{\mathbf{x}}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \bar{\mathbf{x}}_t, \mathbf{c}_t}^{-1} L_{\bar{\mathbf{x}}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}^{-1} \\ &= L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} D_{\bar{\mathbf{x}}_t | \mathbf{x}_{t-1}, \mathbf{c}_t}^{-1} D_{\bar{\mathbf{x}}_t | \bar{\mathbf{x}}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \bar{\mathbf{x}}_t, \mathbf{c}_t}^{-1} L_{\bar{\mathbf{x}}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}^{-1} \\ &\quad L_{\mathbf{x}_{t+1} | \bar{\mathbf{x}}_t, \mathbf{c}_t} D_{\bar{\mathbf{x}}_t | \bar{\mathbf{x}}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \bar{\mathbf{x}}_t, \mathbf{c}_t}^{-1} L_{\bar{\mathbf{x}}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}^{-1}. \end{aligned} \quad (\text{A23})$$

272 where

$$\begin{aligned} D_{\mathbf{x}_t, \bar{\mathbf{x}}_t, \mathbf{x}_{t-1}, \bar{\mathbf{x}}_t, \mathbf{c}_t} h(\mathbf{c}_t) &= D_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t} D_{\bar{\mathbf{x}}_t | \mathbf{x}_{t-1}, \mathbf{c}_t}^{-1} D_{\bar{\mathbf{x}}_t | \bar{\mathbf{x}}_t, \mathbf{c}_t} D_{\mathbf{x}_t | \bar{\mathbf{x}}_t, \mathbf{c}_t}^{-1} h(\mathbf{c}_t) \\ &= \frac{\rho_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t}(\mathbf{x}_t; \mathbf{x}_{t-1}; \mathbf{c}_t) \rho_{\bar{\mathbf{x}}_t | \mathbf{x}_{t-1}, \mathbf{c}_t}(\bar{\mathbf{x}}_t; \mathbf{x}_{t-1}; \mathbf{c}_t)}{\rho_{\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{c}_t}(\mathbf{x}_t; \mathbf{x}_{t-1}; \mathbf{c}_t) \rho_{\bar{\mathbf{x}}_t | \mathbf{x}_{t-1}, \mathbf{c}_t}(\bar{\mathbf{x}}_t; \mathbf{x}_{t-1}; \mathbf{c}_t)} h(\mathbf{c}_t) \\ &\quad k(\mathbf{x}_t; \bar{\mathbf{x}}_t; \mathbf{x}_{t-1}; \bar{\mathbf{x}}_t; \mathbf{c}_t) h(\mathbf{c}_t). \end{aligned} \quad (\text{A24})$$

273 This equation implies that the observed operator \mathbf{MN} on the L.H.S. of Eq. A25 has an inher-
274 ent eigenvalue–eigenfunction decomposition, with the eigenvalues corresponding to the function
275 $k(\mathbf{x}_t; \bar{\mathbf{x}}_t; \mathbf{x}_{t-1}; \bar{\mathbf{x}}_t; \mathbf{c}_t)$ and the eigenfunctions corresponding to the density $\rho_{\bar{\mathbf{x}}_{t+1} | \mathbf{x}_t, \mathbf{c}_t}(\bar{\mathbf{x}}_t; \mathbf{x}_t; \mathbf{c}_t)$.

276 The decomposition in Eq. A25 is similar to the decomposition in nonparametric identification [43, 66].
277 First, Assumption 3 ensures this decomposition is unique. Second, the operator \mathbf{MN} on the L.H.S.
278 has the same spectrum as the diagonal operator $D_{\mathbf{x}_t, \bar{\mathbf{x}}_t, \mathbf{x}_{t-1}, \bar{\mathbf{x}}_t, \mathbf{c}_t}$. Assumption 3 guarantees that the
279 spectrum of the diagonal operator is bounded. Since an operator is bounded by the largest element of
280 its spectrum, Assumption 3 also implies that the operator \mathbf{MN} is bounded, whence we can apply
281 Theorem XV.4.3.5 from [42] to show the uniqueness of the spectral decomposition of bounded linear
282 operators:

$$L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} = C L_{\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{c}_t} P^{-1}; \quad D_{\mathbf{x}_t, \bar{\mathbf{x}}_t, \mathbf{x}_{t-1}, \bar{\mathbf{x}}_t, \mathbf{c}_t} = P D_{\mathbf{x}_t, \bar{\mathbf{x}}_t, \mathbf{x}_{t-1}, \bar{\mathbf{x}}_t, \mathbf{c}_t} P^{-1} \quad (\text{A25})$$

283 where C is a scalar accounting for scaling indeterminacy and P is a permutation on the order of
284 elements in $D_{\bar{\mathbf{x}}_t | \hat{\mathbf{c}}_t}$, as discussed in [42]. These forms of indeterminacy are analogous to those in
285 eigendecomposition, which can be viewed as a finite-dimensional special case.

286 We will show why the uniqueness of spectral decomposition is informative for identifications. First,
287

$$\int_{\bar{\mathbf{x}}_{t+1}} \rho_{\bar{\mathbf{x}}_{t+1} | \bar{\mathbf{x}}_t, \mathbf{c}_t} d\bar{\mathbf{x}}_{t+1} = 1 \quad (\text{A26})$$

287 must hold for every $\hat{\mathbf{c}}_t$ due to normalizing condition, one only solution is to set $C = 1$.

288 Second, Assumption 3 implies that Eq. A25 imply that the eigenvalues $k(\mathbf{x}_t; \bar{\mathbf{x}}_t; \mathbf{x}_{t-1}; \bar{\mathbf{x}}_t; \mathbf{c}_t)$ are
289 distinct for different values \mathbf{c}_t . If several \mathbf{c}_t yield identical eigenvalues, the associated eigenfunctions
290 cannot be uniquely identified, as any linear combination of them remains valid. Therefore, for each
291 \mathbf{x}_t , one can choose $\mathbf{x}_t; \bar{\mathbf{x}}_t; \mathbf{x}_{t-1}; \bar{\mathbf{x}}_t$ such that the eigenvalues differ for all \mathbf{c}_t .

292 Ultimately, the unorder of eigenvalues/eigenfunctions is left. The operator, $L_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}$, corresponding
 293 to the set $\bar{f}p_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}(j \mathbf{x}_t; \mathbf{c}_t)g$ for all $\mathbf{x}_t; \mathbf{c}_t$, admits a unique solution (ordering ambiguity of
 294 eigendecomposition only changes the entry position):

$$\bar{f}p_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}(j \mathbf{x}_t; \mathbf{c}_t)g = \bar{f}p_{\mathbf{x}_{t+1}|\hat{\mathbf{x}}_t, \hat{\mathbf{c}}_t}(\mathbf{x}_{t+1} j \hat{\mathbf{x}}_t; \hat{\mathbf{c}}_t)g; \quad \text{for all } \mathbf{x}_t; \mathbf{c}_t; \hat{\mathbf{x}}_t; \hat{\mathbf{c}}_t \quad (\text{A27})$$

295 Due to the set is unorder, the only way to match the R.H.S. with the L.H.S. in a consistent order is to
 296 exchange the conditioning variables, that is,

$$\begin{aligned} & \bar{f}p_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}(j \mathbf{x}_t^{(1)}; \mathbf{c}_t^{(1)}); p_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}(j \mathbf{x}_t^{(2)}; \mathbf{c}_t^{(2)}); \dots; g \\ & = \bar{f}p_{\mathbf{x}_{t+1}|\hat{\mathbf{x}}_t, \hat{\mathbf{c}}_t}(j \hat{\mathbf{x}}_t^{(1)}; \hat{\mathbf{c}}_t^{(1)}); p_{\mathbf{x}_{t+1}|\hat{\mathbf{x}}_t, \hat{\mathbf{c}}_t}(j \hat{\mathbf{x}}_t^{(2)}; \hat{\mathbf{c}}_t^{(2)}); \dots; g \\ &) \quad [p_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}(j \mathbf{x}_t^{(\pi(1))}; \mathbf{c}_t^{(\pi(1))}); p_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}(j \mathbf{x}_t^{(\pi(2))}; \mathbf{c}_t^{(\pi(2))}); \dots;] \\ & = [p_{\mathbf{x}_{t+1}|\hat{\mathbf{x}}_t, \hat{\mathbf{c}}_t}(j \hat{\mathbf{x}}_t^{(\pi(1))}; \hat{\mathbf{c}}_t^{(\pi(1))}); p_{\mathbf{x}_{t+1}|\hat{\mathbf{x}}_t, \hat{\mathbf{c}}_t}(j \hat{\mathbf{x}}_t^{(\pi(2))}; \hat{\mathbf{c}}_t^{(\pi(2))}); \dots;] \end{aligned} \quad (\text{A28})$$

297 where superscript $()$ denotes the index of the conditioning variables $[\mathbf{x}_t; \mathbf{c}_t]$, and π is reindexing the
 298 conditioning variables. We use a relabeling map H to represent its corresponding value mapping:

$$p_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}(j H(\mathbf{x}_t; \mathbf{c}_t)) = p_{\mathbf{x}_{t+1}|\hat{\mathbf{x}}_t, \hat{\mathbf{c}}_t}(j \hat{\mathbf{x}}_t; \hat{\mathbf{c}}_t); \quad \text{for all } \mathbf{x}_t; \mathbf{c}_t; \hat{\mathbf{x}}_t; \hat{\mathbf{c}}_t \quad (\text{A29})$$

299 By Assumption 3, different \mathbf{c}_t corresponds to different $p_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}(j H(\mathbf{x}_t; \mathbf{c}_t))$, which indicates
 300 that there is no repeated element in $\bar{f}p_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}(j H(\mathbf{x}_t; \mathbf{c}_t))g$ and $\bar{f}p_{\mathbf{x}_{t+1}|\hat{\mathbf{x}}_t, \hat{\mathbf{c}}_t}(j \hat{\mathbf{x}}_t; \hat{\mathbf{c}}_t)g$. Such
 301 uniqueness ensure that the relabelling map H is one-to-one.

302 Furthermore, Assumption 3 implies that $p_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}(j H(\mathbf{x}_t; \mathbf{c}_t))$ corresponds a unique $H(\mathbf{x}_t; \mathbf{c}_t)$.
 303 The same holds for the $p_{\mathbf{x}_{t+1}|\hat{\mathbf{x}}_t, \hat{\mathbf{c}}_t}(j \hat{\mathbf{x}}_t; \hat{\mathbf{c}}_t)$, implying that

$$p_{\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{c}_t}(j H(\mathbf{x}_t; \mathbf{c}_t)) = p_{\mathbf{x}_{t+1}|\hat{\mathbf{x}}_t, \hat{\mathbf{c}}_t}(j \hat{\mathbf{x}}_t; \hat{\mathbf{c}}_t) \Rightarrow \hat{\mathbf{x}}_t; \hat{\mathbf{c}}_t = H(\mathbf{x}_t; \mathbf{c}_t) \quad (\text{A30})$$

304 Since the observation \mathbf{x}_t is known and suppose $\hat{\mathbf{x}}_t = \mathbf{x}_t$, this relationship indeed represents an
 305 invertible transformation between $\hat{\mathbf{c}}_t$ and \mathbf{c}_t as

$$\hat{\mathbf{c}}_t = h(\mathbf{c}_t); \quad (\text{A31})$$

306 which ensures that $p(\mathbf{c}_t | \mathbf{x}_{t-2:t+1})$ can be identifiable up to an invertible transformation on the latent
 307 variables $\hat{\mathbf{c}}_t = h(\mathbf{c}_t)$ \square

308 B.3 Theory-Algorithm Alignment

309 Here, we provide a more detailed description of our theoretical foundations, model design guidance,
 310 and algorithmic implementation. This complements the high-level summary in the main paper and
 311 offers additional context about the technical depth behind our contributions.

312 B.3.1 Non-Parametric Identifiability Theory

313 We establish this non-parametric identifiability result that gives sufficient conditions for recovering
 314 latent contexts from reinforcement learning (RL) trajectories using short temporal blocks. Each block
 315 includes a small number of future steps, which allows the model to reason about both the immediate
 316 past and the near future. Formally, we prove that under mild and broadly applicable assumptions,
 317 the latent context \mathbf{c}_t driving the generative process of the observed states and actions $(\mathbf{x}_t; \mathbf{a}_t)$ can be
 318 recovered up to an equivalence class. This identifiability guarantee is important because it shows that
 319 latent-aware planning can be theoretically justified even when the environment contains unobserved
 320 factors or task-dependent variations.

321 B.3.2 Model Design Guidance

322 The theoretical result directly guides the design of our causal diffusion model. To leverage Theorem 1,
 323 the diffusion model must not only generate trajectories but also recover the true latent factors.
 324 Concretely, the model must:

- 325 1. capture **temporal dependencies** across short blocks of states and actions;
- 326 2. jointly model **observable and latent variables**;

327 3. enforce **conditions for identifiability**, ensuring that the latent \mathbf{c}_t can be isolated from the
 328 observed sequence.

329 These design requirements inform our noise schedule and the coupling of autoregressive denoising
 330 with latent refinement.

331 This provides guidance for the algorithm design:

332 **Autoregressive Denoising.** We model temporal dependencies over both observable and latent
 333 variables using an autoregressive diffusion process. At each step, \mathbf{x}_t is denoised while conditioning
 334 on partially denoised past states and inferred latent variables from a short temporal block (Section 4.1).
 335 This schedule results in a structured temporal-latent modeling process that better preserves long-range
 336 dependencies.

337 **Backward Refinement.** To explicitly identify latent contexts whose posterior depends on future
 338 observations (e.g., \mathbf{x}_{t+1} , guided by Theorem 1), we introduce a backward refinement step. At the
 339 second-to-last denoising stage, we refine a partial state $\mathbf{x}_t^{k_1}$ using the initial estimate of $\hat{\mathbf{c}}_t$ sampled
 340 from the prior and \mathbf{x}_{t+1} as additional evidence. The refined $\hat{\mathbf{c}}_t$ is then used to produce the final
 341 denoised state \mathbf{x}_t^0 . During training, this backward refinement is enforced to satisfy the identifiability
 342 conditions. At inference time (zig-zag sampling), we substitute \mathbf{x}_{t+1} with a predicted estimate to
 343 maintain efficiency.

344 **Unification.** The autoregressive denoising and backward refinement are integrated into a single
 345 noise schedule, enabling joint modeling of temporal dependencies and latent variables. Our imple-
 346 mentation follows a four-step refinement scheme but can be extended to more steps if needed. Notably,
 347 despite the additional refinement, the method remains computationally efficient (see Appendix I.1).
 348 Further acceleration is possible via Picard iteration, which parallelizes refinement steps and reduces
 349 inference runtime by about 25–30%.

350 B.3.3 Discussion on Assumptions

351 B.3.4 Relaxing Assumption 1 (beyond first-order Markov)

352 We can relax the first-order Markov assumption to an n -order Markov structure with delayed/cumu-
 353 lative influences without altering the core identifiability argument. Suppose the generative process
 354 satisfies

$$p(\mathbf{x}_{t+1} \mid \mathbf{x}_{1:t}; \mathbf{a}_{1:t}; \mathbf{c}_{1:t}) = p(\mathbf{x}_{t+1} \mid \mathbf{x}_{t-n+1:t}; \mathbf{a}_{t-n+1:t}; \mathbf{c}_{t-n+1:t})$$

355 and that the conditioning sets across *non overlapping* lags exhibit block-wise conditional inde-
 356 pendence (the same separation conditions used in Theorem 1). Then there exists a finite window
 357 of observations whose statistics identify the contemporaneous block $[\mathbf{c}_t; \mathbf{x}_t]$ up to an invertible
 358 reparameterization.

359 **Concrete identification statement.** Let $W_t = \mathbf{x}_{t-2n:t+2n}$ denote a $4n+1$ -length observation
 360 window.³ Assume: (i) time direction is known (so $[\mathbf{c}_t; \mathbf{x}_t] \neq [\mathbf{c}_{t+1}; \mathbf{x}_{t+1}]$ is oriented); (ii) the
 361 variability (support) conditions from Theorem 1 hold for the n -lag blocks; and (iii) block-wise
 362 independence across non-overlapping lags is satisfied. Then there exists an invertible map H such
 363 that

$$[\mathbf{c}_t; \mathbf{x}_t] = H W_t$$

364 so \mathbf{c}_t (and \mathbf{x}_t) are identifiable up to an invertible transformation from a finite window of observations.

365 **Illustration for $n=2$.** When $n=2$, block-wise separations allow identification of the joint variables
 366 $[\mathbf{c}_t; \mathbf{c}_{t+1}; \mathbf{x}_t; \mathbf{x}_{t+1}]$ from $\mathbf{x}_{t-4:t+3}$ (length 8+1). Knowing the temporal direction disambiguates
 367 $[\mathbf{c}_t; \mathbf{x}_t]$ from $[\mathbf{c}_{t+1}; \mathbf{x}_{t+1}]$. Because \mathbf{x}_t is observed, we obtain $\mathbf{c}_t = h(\mathbf{x}_{t-4:t+3})$ for some invertible
 368 h , and thus the contemporaneous pair $[\mathbf{c}_t; \mathbf{x}_t]$ is identified.

³Any window of length at least $4n+1$ suffices; we state one concrete choice for clarity.

Latent Type	Block size	Probing Acc	R^2
Delayed	6	0.81	0.72
Delayed	8	0.85	0.78
Delayed	10	0.88	0.81
Delayed	20	0.91	0.86
Cumulative	6	0.84	0.75
Cumulative	8	0.87	0.79
Cumulative	10	0.89	0.83
Cumulative	20	0.93	0.88

Table A2: Identification under delayed and cumulative latent effects. Larger is better.

369 **Connection to delayed/cumulative rewards.** Delayed and cumulative effects used in §X fit
370 naturally in the n -order view. For a delay ℓ ,

$$\mathbf{r}_{t+\ell} = \mathbf{x}_{t+\ell} \mathbf{a}_{t+\ell} \mathbf{c}_t \quad (\text{delayed effect}),$$

371 while cumulative influence over a horizon L can be written as

$$\mathbf{r}_{t+k} = \sum_{i=0}^k \mathbf{x}_{t+i} \mathbf{a}_{t+i} \mathbf{c}_t \quad ; \quad k = 0, \dots, L-1;$$

372 both of which are encompassed by the n -order Markov factorization above. Our cheetah variants
373 instantiate these with, e.g., $r_{t+\ell} = \mathbf{v}_{t+\ell} \mathbf{c}_t k^2$ (delayed) and $r_{t+k} = \sum_{i=0}^k \mathbf{v}_{t+i} \mathbf{c}_t k^2$ (cumulative),
374 where \mathbf{v}_t denotes speed; the identification results remain valid.

375 **Results.** We evaluate identification under delayed and cumulative latent effects in the Cheetah
376 environment using observation windows of length 6; 8; 10; and 20. In all cases, linear probes recover
377 the latent with high accuracy, and performance improves monotonically with longer context. For
378 *delayed* effects, probing accuracy rises from 0.81 to 0.91 and R^2 from 0.72 to 0.86 as block size
379 increases from 6 to 20. For *cumulative* effects, probing accuracy increases from 0.84 to 0.93 and R^2
380 from 0.75 to 0.88 over the same range. These results confirm that (i) the latent \mathbf{c}_t is behaviorally
381 consequential in non-first-order settings and (ii) moderate temporal context suffices for accurate
382 recovery, supporting our relaxed n -order Markov analysis.

383 B.3.5 Cases in Assumption 2.

384 The assumption of the injectivity of a linear operator is commonly employed in the nonparametric
385 identification [43, 66, 44]. Intuitively, it means that different input distributions of a linear operator
386 correspond to different output distributions of that operator. For a better understanding, we provide
387 several examples in Fu et al. [45] that describe the mapping from p_a to p_b , where a and b are random
388 variables:

389 **Example 1** (Invertible). $b = g(a)$, where g is an invertible function.

390 **Example 2** (Additive). $b = a + \epsilon$, where $p(\epsilon)$ must not vanish everywhere after the Fourier transform.

391 **Example 3** (Nonlinear Additive). $b = g(a) + \epsilon$, where conditions from **Examples 1-2** are required.

392 **Example 4** (Post-nonlinear). $b = g_1(g_2(a) + \epsilon)$, a post-nonlinear model with invertible nonlinear
393 functions g_1, g_2 , combining the assumptions in **Examples 1-3**.

394 **Example 5** (Nonlinear with Exponential Family). $b = g(a; \theta)$, where the joint distribution $p(a; b)$
395 follows an exponential family.

396 **Example 6** (Nonparametric). $b = g(a; \theta)$, a general nonlinear formulation. Certain deviations from
397 the nonlinear additive model (**Example 3**), e.g., polynomial perturbations, can still be tractable.

398 B.4 ELBO

399 In this section, we provide analysis on the \mathbf{x}^0 -prediction Mean Squared Error (MSE) loss objectives
400 used in the Denoise-and-Refine Mechanism of Ada-Diffuser. Our main argument establishes
401 that minimizing the reconstruction losses L_{prior} and L_{post} corresponds to optimizing an ELBO on the

402 conditional log-likelihood of the clean observation \mathbf{x}_t^0 , given a noisy observation \mathbf{x}_t^k and an inferred
 403 latent context \mathbf{c}_t .

404 Let $\mathbf{x}_t^0 \sim q(\mathbf{x}_t^0)$ be a clean data sample from the true data distribution at sequence time step t . Let \mathbf{c}_t
 405 be the inferred latent context relevant to \mathbf{x}_t^0 .

The forward diffusion process gradually adds Gaussian noise to \mathbf{x}_t^0 over K diffusion steps:

$$q(\mathbf{x}_t^k | \mathbf{x}_t^{k-1}) = \mathcal{N}(\mathbf{x}_t^k; \rho_{k-1} \mathbf{x}_t^{k-1}; (1 - \rho_k) \mathbf{I})$$

for $k \geq 1; \dots; K$, where $\rho_k \in (0; 1)$ are predefined noise schedule parameters. This process allows
 sampling \mathbf{x}_t^k directly from \mathbf{x}_t^0 :

$$\mathbf{x}_t^k = \rho_{k-1} \mathbf{x}_t^0 + \sqrt{1 - \rho_k} \boldsymbol{\epsilon}; \quad \text{where } \boldsymbol{\epsilon} \sim \mathcal{N}(0; \mathbf{I}), \text{ and } \rho_k = \prod_{i=1}^k \rho_i$$

406 The reverse process $p_\theta(\mathbf{x}_t^{k-1} | \mathbf{x}_t^k; \mathbf{c}_t)$ that parameterized by θ aims to denoise \mathbf{x}_t^k to \mathbf{x}_t^{k-1} conditioned
 407 on \mathbf{c}_t .

The derivation of the ELBO for diffusion models is standard following DDPM related derivations [31,
 22]. The conditional log-likelihood $\log p_\theta(\mathbf{x}_t^0 | \mathbf{c}_t)$ can be lower-bounded using the ELBO:

$$\log p_\theta(\mathbf{x}_t^0 | \mathbf{c}_t) \geq \mathbb{E}_{q(\mathbf{x}_t^{1:K} | \mathbf{x}_t^0)} \left[\log p_\theta(\mathbf{x}_t^K | \mathbf{c}_t) + \sum_{k=1}^{K-1} \log \frac{p_\theta(\mathbf{x}_t^{k-1} | \mathbf{x}_t^k; \mathbf{c}_t)}{q(\mathbf{x}_t^{k-1} | \mathbf{x}_t^k; \mathbf{x}_t^0)} \right]$$

408 Assuming p_θ satisfies Markov Property (i.e., $p_\theta(\mathbf{x}_t^{k-1} | \mathbf{x}_t^k; \dots; \mathbf{x}_t^K; \mathbf{c}_t) = p_\theta(\mathbf{x}_t^{k-1} | \mathbf{x}_t^k; \mathbf{c}_t)$), which
 409 is a standard structural assumption for diffusion models, the ELBO can be rewritten as:

$$\begin{aligned} \log p_\theta(\mathbf{x}_t^0 | \mathbf{c}_t) &\geq \underbrace{\mathbb{E}_{q(\mathbf{x}_t^1 | \mathbf{x}_t^0)} [\log p_\theta(\mathbf{x}_t^1 | \mathbf{x}_t^0; \mathbf{c}_t)]}_{L_0} \\ &\quad - \sum_{k=2}^K \underbrace{\mathbb{E}_{q(\mathbf{x}_t^k | \mathbf{x}_t^0)} [D_{KL}(q(\mathbf{x}_t^{k-1} | \mathbf{x}_t^k; \mathbf{x}_t^0) || p_\theta(\mathbf{x}_t^{k-1} | \mathbf{x}_t^k; \mathbf{c}_t))]}_{L_{k-1}} \\ &\quad - \underbrace{D_{KL}(q(\mathbf{x}_t^K | \mathbf{x}_t^0) || p_\theta(\mathbf{x}_t^K | \mathbf{c}_t))}_{L_K}; \end{aligned}$$

410 This inequality holds with equality if and only if the model's true posterior over the latent diffusion
 411 path, $p_\theta(\mathbf{x}_t^{1:K} | \mathbf{x}_t^0; \mathbf{c}_t)$, is identical to the approximate posterior used to derive the ELBO, which is
 412 the forward noising process $q(\mathbf{x}_t^{1:K} | \mathbf{x}_t^0)$. This bound can also include an additive constant $C(\mathbf{x}_t^0; \mathbf{c}_t)$
 413 which does not depend on the model parameters and is thus typically omitted when focusing on
 414 terms relevant to parameter optimization.

415 To maximize $\log p_\theta(\mathbf{x}_t^0 | \mathbf{c}_t)$, we aim to maximize this lower bound by optimizing L_0 (i.e., maximizing
 416 this term) and each L_{k-1} term (i.e., minimizing these D_{KL} terms, as they appear with a negative
 417 sign). The term L_K is often treated as a constant (or absorbed into $C(\mathbf{x}_t^0; \mathbf{c}_t)$) if $p_\theta(\mathbf{x}_t^K | \mathbf{c}_t)$ is set to a
 418 standard Gaussian $\mathcal{N}(0; \mathbf{I})$ and $\rho_K = 0$.

We parameterize the reverse process $p_\theta(\mathbf{x}_t^{k-1} | \mathbf{x}_t^k; \mathbf{c}_t)$ as a Gaussian:

$$p_\theta(\mathbf{x}_t^{k-1} | \mathbf{x}_t^k; \mathbf{c}_t) = \mathcal{N}(\mathbf{x}_t^{k-1}; \theta(\mathbf{x}_t^k; k; \mathbf{c}_t); \frac{2}{k} \mathbf{I})$$

The true posterior step $q(\mathbf{x}_t^{k-1} | \mathbf{x}_t^k; \mathbf{x}_t^0)$ is also Gaussian:

$$q(\mathbf{x}_t^{k-1} | \mathbf{x}_t^k; \mathbf{x}_t^0) = \mathcal{N}(\mathbf{x}_t^{k-1}; \tilde{\mu}_k(\mathbf{x}_t^k; \mathbf{x}_t^0); \tilde{\Sigma}_k \mathbf{I})$$

419 where $\tilde{\mu}_k(\mathbf{x}_t^k; \mathbf{x}_t^0) = \frac{\sqrt{\bar{\alpha}_k} \bar{\alpha}_k (1 - \alpha_k)}{1 - \bar{\alpha}_k} \mathbf{x}_t^0 + \frac{\sqrt{\bar{\alpha}_k} (1 - \bar{\alpha}_k)}{1 - \bar{\alpha}_k} \mathbf{x}_t^k$ and $\tilde{\Sigma}_k = \frac{1 - \bar{\alpha}_k}{1 - \bar{\alpha}_k} (1 - \rho_k)$ is the variance.

For an \mathbf{x}^0 -prediction model, denoted as $\theta(\mathbf{x}_t^k; k; \mathbf{c}_t)$ in the main paper, that aims to predict \mathbf{x}_t^0 from
 the noisy input \mathbf{x}_t^k and context \mathbf{c}_t , the mean of the reverse model θ can be expressed as:

$$\theta(\mathbf{x}_t^k; k; \mathbf{c}_t) = \frac{\rho_{k-1}}{1 - \rho_k} \theta(\mathbf{x}_t^{k-1}; k; \mathbf{c}_t) + \frac{\rho_k}{1 - \rho_k} \mathbf{x}_t^k$$

420 Choosing $\frac{2}{k} = -\frac{2}{k}$, the KL divergence term L_{k-1} simplifies to:

$$\begin{aligned} L_{k-1} &= \mathbb{E}_{q(\mathbf{x}_t^k | \mathbf{x}_t^0)} \frac{1}{2} \frac{1}{k} \|\mathbf{x}_t^k - \mathbf{x}_t^0\|_{\theta(\mathbf{x}_t^k; k; \mathbf{c}_t)}^2 + C'_k \\ &= \mathbb{E}_{\mathbf{x}_t^0} \frac{1}{2} \frac{1}{k} \frac{\rho_{k-1}(1-k)}{1-k} \|\mathbf{x}_t^0 - \theta(\rho_{k-1} \mathbf{x}_t^0 + \rho_{1-k}; k; \mathbf{c}_t)\|_{\theta}^2 + C'_k \end{aligned}$$

421 where C'_k are constants not depending on θ . The expectation $\mathbb{E}_{\mathbf{x}_t^0}$, denotes averaging over clean data
422 \mathbf{x}_t^0 and the noise ϵ used to construct \mathbf{x}_t^k . Thus, maximizing the ELBO contribution from L_{k-1} is
423 equivalent to minimizing the following weighted MSE term:

$$\mathbb{E}_{\mathbf{x}_t^0, \mathbf{c}_t} w(k) \|\mathbf{x}_t^0 - \theta(\rho_{k-1} \mathbf{x}_t^0 + \rho_{1-k}; k; \mathbf{c}_t)\|_{\theta}^2 \quad (\text{A32})$$

424 where $w(k) = \frac{1}{2\sigma_k^2} \frac{\sqrt{\alpha_k} (1-\alpha_k)}{1-\alpha_k}$ is a positive weighting factor.

The term $L_0 = \mathbb{E}_{q(\mathbf{x}_t^1 | \mathbf{x}_t^0)} [\log p_{\theta}(\mathbf{x}_t^0 | \mathbf{x}_t^1; \mathbf{c}_t)]$ can also be made proportional to an MSE if $p_{\theta}(\mathbf{x}_t^0 | \mathbf{x}_t^1; \mathbf{c}_t)$ is a Gaussian centered at $\theta(\mathbf{x}_t^1; 1; \mathbf{c}_t)$:

$$\log p_{\theta}(\mathbf{x}_t^0 | \mathbf{x}_t^1; \mathbf{c}_t) = -\frac{1}{2} \|\mathbf{x}_t^0 - \theta(\mathbf{x}_t^1; 1; \mathbf{c}_t)\|_{\theta}^2 + \text{const}$$

425 Maximizing L_0 is then equivalent to minimizing this MSE.

The diffusion model θ is typically trained by minimizing a simplified objective (e.g., [31]), often an unweighted or equally weighted sum of these MSE terms over uniformly sampled diffusion steps $k \in [1; K]$ and data \mathbf{x}_t^0 :

$$L_{\text{simple}}(\theta) = \mathbb{E}_{k \sim U[1, K], \mathbf{x}_t^0, \mathbf{c}_t} \|\mathbf{x}_t^0 - \theta(\rho_{k-1} \mathbf{x}_t^0 + \rho_{1-k}; k; \mathbf{c}_t)\|_{\theta}^2$$

426 This simplification is justified by arguing that reweighting terms $w(k)$ in Equation A32 can be
427 absorbed into the network or do not significantly alter the optimal solution for expressive models,
428 allowing $w(k)$ to be effectively set to 1.

The Denoise-and-Refine losses are:

$$L_{\text{prior}} = \mathbb{E}_{\mathbf{x}_t^0, \mathbf{c}_t^{\text{prior}}} \|\mathbf{x}_t^0 - \theta(\mathbf{x}_t^{k_i}; k_i; \mathbf{c}_t^{\text{prior}})\|_{\theta}^2$$

$$L_{\text{post}} = \mathbb{E}_{\mathbf{x}_t^0, \mathbf{c}_t^{\text{post}}} \|\mathbf{x}_t^0 - \theta(\mathbf{x}_t^{k_i}; k_i; \mathbf{c}_t^{\text{post}})\|_{\theta}^2$$

429 where $\mathbf{x}_t^{k_i} = \rho_{k_i-1} \mathbf{x}_t^0 + \rho_{1-k_i}$, and k_i is the specific input noise level for the observation \mathbf{x}_t
430 determined by the causal denoising schedule $k_i = \frac{i}{T} K$. These losses, L_{prior} and L_{post} , are specific
431 instances of the simplified MSE loss objective in equation A32 with $w(k_i) = 1$, conditioned on the
432 inferred contexts $\mathbf{c}_t^{\text{prior}}$ and $\mathbf{c}_t^{\text{post}}$ respectively. Consequently, minimizing these MSE losses directly
433 optimizes the corresponding terms in the ELBO for $\log p_{\theta}(\mathbf{x}_t^0 | \mathbf{c}_t)$.

434 Therefore, we have proven that minimizing L_{prior} and L_{post} as defined in the Denoise-and-Refine mech-
435 anism serves to maximize a variational lower bound on the conditional log-likelihood $\log p_{\theta}(\mathbf{x}_t^0 | \mathbf{c}_t)$.
436 The underlying diffusion model $\theta(\cdot; k; \cdot)$ is trained to be proficient at denoising from a range of
437 noise levels k , as captured by objectives such as L_{simple} . The specific monotonically increasing
438 noise schedule k_i used in L_{prior} and L_{post} represents a particular instance from this range of noise
439 levels. Thus, these objectives are theoretically grounded in the principles of variational inference for
440 diffusion models, adapted to conditioning on the inferred latent context \mathbf{c}_t and applied at specific
441 noise levels relevant to the autoregressive denoising process of Ada-Diffuser.

442 C Summary on Different MDPs

443 Our work considers a contextual POMDP setting with an evolving latent process, which naturally
444 relates to several established MDP formulations, including contextual MDPs [67], hidden-parameter
445 MDPs (HiP-MDPs) [37], and their variants. In this section, we provide formal definitions of these
446 models and discuss their relationships and distinctions.

447 **C.1 Contextual MDPs**

448 A contextual Markov decision process (CMDP) [67] is defined by the tuple $\langle \mathcal{C}; S; A; M \rangle$, where \mathcal{C} is
 449 the context space, S is the state space, and A is the action space. The mapping M assigns to each
 450 context $c \in \mathcal{C}$ a set of MDP parameters $M(c) = \langle R^c; T^c \rangle$, where R^c and T^c are the reward and
 451 transition functions associated with context c .

452 Sodhani et al. [68] and Liang et al. [20] extend the CMDP framework to settings in which the context
 453 variable \mathbf{c} evolves according to its own Markovian dynamics $p(\mathbf{c}_{t+1} | \mathbf{c}_t)$, closely aligning with our
 454 formulation of a latent process evolving over time.

455 **C.2 Hidden-Parameter MDPs**

456 Hidden-Parameter MDPs (HiP-MDPs) [37] are defined by the tuple $M = \langle S; A; \Theta; T; R; \gamma; P_\Theta \rangle$,
 457 where S is the state space, A is the action space, and Θ is the space of task-specific latent parameters.
 458 For each $\theta \in \Theta$, the transition and reward functions are given by $T_\theta : S \times A \rightarrow \mathcal{P}(S)$ and
 459 $R_\theta : S \times A \rightarrow \mathbb{R}$, respectively. The parameter θ is sampled from a prior distribution P_Θ at the
 460 beginning of an episode and remains fixed during the episode. The discount factor is denoted by
 461 $\gamma \in [0, 1)$. This framework defines a family of MDPs indexed by the latent parameter θ , with each
 462 θ inducing a different set of dynamics and reward functions. It can be seen as a special case of a
 463 contextual MDP where the context is latent and fixed per episode. Xie et al. [14] further generalize
 464 this framework by allowing the task parameter θ to evolve dynamically across episodes, rather than
 465 being fixed.

466 Bayes-Adaptive MDPs (BAMDPs) are closely related to both HiP-MDPs and contextual MDPs
 467 (CMDPs). In BAMDPs, the agent maintains a posterior distribution over MDPs based on its
 468 interaction history. Specifically, it maintains a belief $b_t(R; T) = p(R; T | \tau_{:t})$, where $\tau_{:t} =$
 469 $\langle \mathbf{s}_0; \mathbf{a}_0; r_0; \dots; \mathbf{s}_t \rangle$ denotes the trajectory observed up to time t . This belief captures the agent’s
 470 uncertainty about the underlying transition and reward functions.

471 The transition and reward functions can then be defined in expectation over this posterior, effectively
 472 conditioning decision-making on the belief b_t . When the environment is driven by hidden contextual
 473 variables or latent task parameters, such as in CMDPs or HiP-MDPs—this belief can be interpreted as
 474 a distribution over these latent variables. In this view, BAMDPs provide a non-parametric framework
 475 for reasoning over hidden structure, while approaches like ours explicitly model such latent variables
 476 and infer their posterior distributions using amortized inference. Both aim to enable adaptive planning
 477 and learning under uncertainty, but differ in how latent structure is represented and inferred.

478 **C.3 Discussions and Comparisons**

479 The key distinction between contextual MDPs and hidden-parameter MDPs lies in how the latent
 480 factors are represented: contextual MDPs explicitly treat them as latent variables, while HiP-MDPs
 481 model them implicitly as parameters governing the transition and reward functions. In our work, we
 482 adopt the contextual MDP perspective, where the latent process is modeled as a random variable that
 483 evolves over time.

484 However, our identification theory, focused on recovering the posterior distribution over latent
 485 variables, also applies to the HiP-MDP setting. Once the posterior over the hidden parameters is
 486 identified, the corresponding transition and reward functions can be recovered as well.

487 Additionally, our framework, which models a factorization over observed states and latent variables,
 488 is conceptually related to factored MDPs [41]. In a factored MDP, the state space S is represented as
 489 a set of variables $S = \langle s^{(1)}; s^{(2)}; \dots; s^{(n)} \rangle$, and the transition and reward functions are decomposed
 490 over these factors:

$$T(s' | s; \mathbf{a}) = \prod_{i=1}^n T_i(s^{(i)} | \text{Pa}_T^{(i)}(s; \mathbf{a})) ; \quad R(s; \mathbf{a}) = \prod_{j=1}^n R_j(\text{Pa}_R^{(j)}(s; \mathbf{a})) ;$$

491 where $\text{Pa}_T^{(i)}$ and $\text{Pa}_R^{(j)}$ denote the parent variables (i.e., dependencies) for each transition and reward
 492 component, respectively. Our framework, while not relying on an explicit graphical structure, shares
 493 conceptual similarities with factored MDPs [41] through its coarse-grained factorization over observed
 494 states and latent variables. Specifically, we distinguish between latent variables that affect the

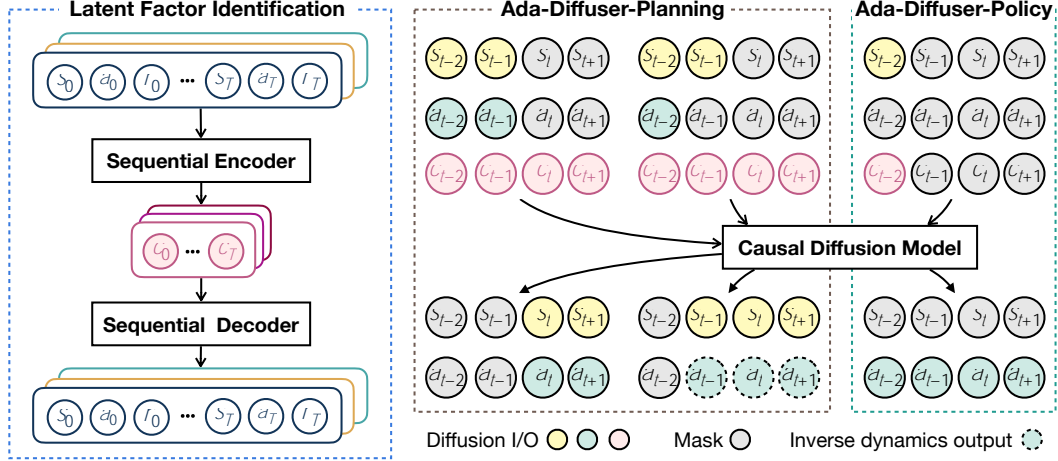


Figure A1: **Overview of the Ada-Diffuser framework.** The modular design consists of two main stages: latent context identification (Stage 1, Section 4.2), followed by a causal diffusion process (Stage 2, Section 4.3) that models the generative structure of the trajectories. The learned model is then used for planning or policy learning conditioned on the inferred latent context.

495 transition dynamics and those that affect the reward function. Formally, we express the generative
 496 process as:

$$T(s_{t+1} | s_t; a_t; c_t^s); \quad R(r_t | s_t; a_t; c_t^r);$$

497 where c_t^s and c_t^r are distinct (or potentially overlapping) latent factors that influence transitions and
 498 rewards, respectively. This separation enables flexible modeling of partially observable environments
 499 where different unobserved processes govern the dynamics and task objectives.

500 D Details on Ada-Diffuser




Figure A2: An illustration of the zig-zag sampling process with a block of 4 time steps. # and j indicate denoising and identity mapping, respectively.

501 **D.1 Full Algorithm and Results**

502 As illustrated in Fig. A1, our framework consists of two stages: latent factor identification and
 503 diffusion-based planning or policy learning. Below, we provide the algorithmic pseudocode for both
 504 stages. Specifically, Algorithm A1 describes Stage 1: latent factor identification, while Algorithms A2
 505 and A3 correspond to Ada-Diffuser-Planner and Ada-Diffuser-Policy, respectively.

506 For clarity, we omit the detailed step-by-step procedures for denoise-and-refine and zig-zag sampling
 507 (Lines 7–8, 11, and 19–22 in Algorithm A2; Lines 6–7 and 13 in Algorithm A3), as these are fully
 508 described in Section 4.3. For Ada-Diffuser-Policy, we show a Diffusion Policy (DP)-based
 509 algorithm, which provides a general framework for multi-step action generation. In the IDQL-based
 510 variant, both the action execution horizon and observation horizon are set to 1, corresponding to
 511 single-step policy inference conditioned only on the current observation.

Algorithm A1: Latent Factor Identification.

- 1: **Input:** offline dataset D
 - 2: Randomly initialize decoder $p_\theta(\mathbf{s}_{t+1}; r_t | \mathbf{s}; \mathbf{a}; \mathbf{c})$,
 encoder $q_\psi(\mathbf{c}_t | \mathbf{s}_{t-T_x:t+1}; \mathbf{a}_{t-T_x:t+1}; r_{t-T_x:t+1})$ and prior network $p_\phi(\mathbf{c}_t | \mathbf{c}_{t-1})$,
 - 3: **while** not done **do**
 - 4: Sample batches of trajectories from D
 - 5: Compute ELBO and update ; ;
 - 6: **end while**
-

512 Additionally, we provide the full results for all experiments: Table A3 reports results for the action-
 513 free setting; Tables A4 and A5 present results for environments with latent factors affecting dynamics
 514 and rewards; and Tables A6, A7, A8, and A9 summarize results for environments without explicitly
 515 modeled latent factors.

516 **D.2 Architecture Choices and Hyper-parameters**

517 We detail the architectural design choices and hyperparameter settings used for model components,
 518 loss functions, and training procedures across all Ada-Diffuser variants under different environ-
 519 ments and benchmarks.

520 **D.2.1 Latent Factor Identification**

521 **Architectures** We use a variational autoencoder (VAE) [48] to optimize the evidence lower bound
 522 (ELBO). The same architectural design is used across all variants of Ada-Diffuser and all
 523 benchmark settings.

524 For the encoder, we first embed states, actions, and rewards using separate MLPs with ReLU
 525 activations. The resulting embeddings are concatenated and passed through a two-layer MLP (each
 526 layer of size 64) followed by a GRU. The GRU output is used to parameterize a Gaussian distribution
 527 from which the latent variables are sampled.

528 The state and reward decoders are implemented as separate MLPs, each consisting of two fully
 529 connected layers of size 64 with ReLU activations. For the prior network, we use the output of the
 530 previous step’s latent distribution embedding (shared GRU) and feed it into a two-layer MLP (each
 531 layer of size 32) to predict the parameters of the prior distribution.

532 **Loss Function** At each time step t , we optimize the following losses:

$$\mathcal{L}_{\text{ELBO}, t} = \underbrace{\mathbb{E}_{q(\mathbf{c}_t | \mathbf{x}_t, \tau_{\mathbf{x}:t+1})} \left[\log p_\theta(\mathbf{x}_t | \mathbf{x}_{t-1}; \mathbf{c}_t) \right]}_{\text{Reconstruction loss}} + \underbrace{\mathbb{D}_{\text{KL}}(q_\psi(\mathbf{c}_t | \mathbf{x}_{t-T_x:t+1}) \parallel p_\phi(\mathbf{c}_t | \mathbf{c}_{t-1}))}_{\text{KL regularization}} :$$

533 Here, \mathbf{x}_t may include different components depending on the setting (e.g., $\mathbf{x}_t = \mathcal{F}\mathbf{s}_t; \mathbf{a}_t \mathcal{G}$ or $\mathbf{x}_t = \mathbf{s}_t$),
 534 and \mathbf{c}_t denotes the latent context variable inferred from a temporal block of observations. The first
 535 term encourages accurate reconstruction of the current observation \mathbf{x}_t conditioned on its immediate

Algorithm A2: Ada-Diffuser-Planner.

```
1: Input: Env, offline dataset  $D$ , pre-trained encoder  $q_\psi$  and prior network  $p_\phi$ 
   observation horizon  $T_o$ , planning horizon  $T_p$ , action execution horizon  $T_a$ , condition  $\mathbf{y}$ 
   // Training
2: Initialize noise predictor  $\theta$ , inverse dynamics model  $f_\phi$ 
3: while not done do
4:   Sample  $\mathbf{x}_{t-T_o:t+T_p}$  from  $D$ 
5:   Sample  $\mathbf{c}_{t:t+T_p}^{\text{prior}}$  and  $\mathbf{c}_{t:t+T_p-2}^{\text{post}}$  from  $p_\phi$  and  $q_\psi$ 
6:   if using inverse dynamics model then
7:     Train Causal Diffusion Model (noise predictor  $\theta$ ) with  $\mathbf{x}_{t-T_o:t}$ ,  $\mathbf{c}_{t-T_o:t}^{\text{prior}}$ , and  $\mathbf{c}_{t-T_o:t}^{\text{post}}$  and
       other conditions  $\mathbf{y}$ , target outputs are  $\mathbf{S}_{t+1:t+T_p}$ 
8:     Train encoder  $q_\psi$  with the contrastive improvement loss  $L_{\text{contrast}}$ 
9:     Train Inverse Dynamics Model  $f_\phi$  to generate actions  $\mathbf{a}_{t+1:t+T_p}$ 
10:  else
11:    Train Causal Diffusion Model (noise predictor  $\theta$ ) with  $\mathbf{x}_{t-T_o:t}$ ,  $\mathbf{c}_{t-T_o:t}^{\text{prior}}$ , and  $\mathbf{c}_{t-T_o:t}^{\text{post}}$  and
      other conditions  $\mathbf{y}$ , target outputs are  $f\mathbf{S}_{t+1:t+T_p}; \mathbf{a}_{t+1:t+T_p}\mathcal{G}$ 
12:    Train encoder  $q_\psi$  with the contrastive improvement loss  $L_{\text{contrast}}$ 
13:  end if
14: end while
   // Execution
15: Initialize environment:  $\mathbf{s}_0 = \text{Env.reset}()$ , set  $t = 0$ 
16: while not done do
   // Observe and infer latent factors
17:   Observe recent trajectory  $\mathbf{X}_{t-T_o:t}$ 
18:   Sample latent variables  $\mathbf{c}_{t:t+T_p}^{\text{prior}}$  from  $p_\phi$ 
   // Generate candidate trajectory
19:   if using inverse dynamics model then
20:     Generate future states (zig-zag sampling)  $\hat{\mathbf{S}}_{t+1:t+T_p}$  conditioned on  $\mathbf{x}_{t-T_o:t}$ ,  $\mathbf{c}_{t:t+T_p}^{\text{prior}}$ , and  $\mathbf{y}$ 
       via learned noise predictor  $\theta$ 
21:     Generate actions  $\hat{\mathbf{a}}_{t+1:t+T_p} = f_\phi(\hat{\mathbf{S}}_{t+1:t+T_p}; \hat{\mathbf{S}}_{t:t+T_p-1})$ 
22:   else
23:     Generate future trajectory  $f\hat{\mathbf{S}}_{t+1:t+T_p}; \hat{\mathbf{a}}_{t+1:t+T_p}\mathcal{G}$  conditioned on  $\mathbf{x}_{t-T_o:t}$ ,  $\mathbf{c}_{t:t+T_p}^{\text{prior}}$ , and  $\mathbf{y}$ 
       via learned noise predictor  $\theta$ 
24:   end if
   // Execute action(s) in environment
25:   for each step  $i = 1$  to  $T_a$  do
26:     Execute  $\hat{\mathbf{a}}_{t+i}$  in Env, observe  $\mathbf{s}_{t+i+1}; r_{t+i}$ 
27:     Append  $(\mathbf{s}_{t+i}; \hat{\mathbf{a}}_{t+i}; r_{t+i})$  to trajectory buffer
28:   end for
29:   Update  $t = t + T_a$ 
30: end while
```

536 past and the latent \mathbf{c}_t , while the second term regularizes the posterior to remain close to the learned
537 prior $p_\phi(\mathbf{c}_t | \mathbf{c}_{t-1})$.

538 We implement the ELBO loss as a weighted combination of the reconstruction loss and the KL
539 divergence:

$$L_{\text{ELBO}} = \sum_{t=1}^{\mathbb{T}-2} \kappa \hat{\mathbf{x}}_t - \mathbf{x}_t \kappa_2^2 + \kappa_{\text{KL}} D_{\text{KL}}(q_\psi(\mathbf{c}_t | \mathbf{x}_{t-T_o:t+1}) | p_\phi(\mathbf{c}_t | \mathbf{c}_{t-1})) ;$$

540 where $\hat{\mathbf{x}}_t$ is the model's reconstruction of the observation \mathbf{x}_t , and κ_{KL} is weighting coefficient. The
541 reconstruction is computed using mean squared error (MSE), and the KL divergence is computed
542 in closed form for Gaussian posteriors and priors. The hyperparameter κ_{KL} is set to be 0.1 and the
543 learning rate is set to be $3e^{-4}$.

Algorithm A3: Ada-Diffuser-Policy (DP-based)

1: **Input:** Env, offline dataset D , pre-trained encoder q_ψ and prior network ρ_ϕ
observation horizon T_o , action generation horizon T_p , action execution horizon T_a , condition \mathbf{y}
// Training
2: Initialize noise predictor θ
3: **while** not done **do**
4: Sample $\mathbf{x}_{t-T_o:t+T_p}$ from D
5: Sample latent variables $\hat{\mathbf{c}}_{t:t+T_p}^{\text{prior}}$, ρ_ϕ , $\hat{\mathbf{c}}_{t:t+T_p-2}^{\text{post}}$, q_ψ
6: Train causal diffusion model (noise predictor θ) to generate actions $\mathbf{a}_{t+1:t+T_p}$, conditioned
on $\mathbf{x}_{t-T_o:t}$, $\hat{\mathbf{c}}_{t:t+T_p}^{\text{prior}}$, $\hat{\mathbf{c}}_{t:t+T_p-2}^{\text{post}}$, and \mathbf{y}
7: Train encoder q_ψ with the contrastive improvement loss L_{contrast}
8: **end while**
// Execution
9: Initialize environment: S_0 , `Env.reset()`, set $t = 0$
10: **while** not done **do**
 // Observe and infer latent factors
11: Observe recent trajectory $\mathbf{x}_{t-T_o:t}$
12: Sample latent variables $\hat{\mathbf{c}}_{t:t+T_p}^{\text{prior}}$, ρ_ϕ
 // Generate actions using causal diffusion model
13: Generate actions (zig-zag sampling) $\hat{\mathbf{a}}_{t+1:t+T_p}$ conditioned on $\mathbf{x}_{t-T_o:t}$, $\hat{\mathbf{c}}_{t:t+T_p}$, and \mathbf{y} via
learned noise predictor θ
 // Execute action(s) in environment
14: **for** each step $i = 1$ to T_a **do**
15: Execute $\hat{\mathbf{a}}_{t+i}$ in Env, observe $S_{t+i+1}; r_{t+i}$
16: Append $(S_{t+i}; \hat{\mathbf{a}}_{t+i}; r_{t+i})$ to trajectory buffer
17: **end for**
18: Update $t = t + T_a$
19: **end while**

Environment	LDP (AF)	Ours (AF)	LDP (AF, SubOpt)	Ours (AF, SubOpt)
Lift	0.67 \pm 0.01	0.78 \pm 0.05	1.00 \pm 0.00	0.98 \pm 0.02
Can	0.78 \pm 0.04	0.85 \pm 0.07	0.98 \pm 0.00	0.98 \pm 0.02
Square	0.47 \pm 0.03	0.54 \pm 0.05	0.83 \pm 0.01	0.89 \pm 0.03

Table A3: **Results (success rate) on action-free demonstrations.** Here, AF and SubOpt indicate using Action-free and suboptimal demonstrations on Robomimic tasks, respectively (following the settings in LDP [54]).

544 D.2.2 Planner

545 For the planner, we consider two scenarios: (i) generating both states and actions, and (ii) generating
546 states only. For the former, we build upon the Diffuser framework [5], which directly models full
547 trajectories. For the latter, we adopt the Decision Diffuser (DD) framework [55], where the model
548 generates future states and uses an inverse dynamics model to recover the corresponding actions via
549 inverse dynamics model.

550 For type (i) (full state-action trajectory generation), we apply our method to the Cheetah and Ant
551 environments. For the noise predictor, we use a 1D U-Net [58] with a kernel size of 5, channel
552 multipliers set to (1, 2, 2, 2), and a base channel width of 32. The model is trained using the Adam
553 optimizer [69] with a learning rate of 3×10^{-4} , a batch size of 64, and for 1 million training steps.
554 We adopt classifier guidance (CG) [31] with gradient guidance on computed return, with a guidance
555 scale $\lambda = 1.5$. The observation horizon is set to 10 for both environments. The planning horizon T_p
556 is set to 16 for Cheetah and 32 for Ant, with an action execution horizon of 1. These hyperparameters
557 are kept consistent across baselines, including Diffuser, DF, MetaDiffuser, and Diffuser combined

558 with LILAC and DynaMITE for the Cheetah and Ant experiments (those in Table 1 and Appendix
559 Table A4). For other components (e.g., VAE) in LDCQ, we employ all the hyperparameters in their
560 original implementation [56].

561 For type (ii) (state-only generation with inverse dynamics), we use a Transformer-based noise
562 predictor with a hidden dimension of 256 and a head dimension of 32. The architecture includes 2
563 DiT blocks for Walker, Kitchen, and Maze2D, and 8 DiT blocks for LIBERO. The model is trained
564 using the Adam optimizer [69] with a learning rate of $3 \cdot 10^{-4}$, a batch size of 128, and for 1 million
565 training steps. The number of diffusion timesteps is 500. The observation horizon is set to 4 for
566 Kitchen, 2 for LIBERO, and 10 for the other environments. The planning horizon T_p is set to 16 for
567 Kitchen, 10 for LIBERO, and 32 for the others. The action execution horizon is 8 for both Kitchen
568 and LIBERO, and 10 for the remaining environments. For the inverse dynamics model, we use an
569 MLP-based diffusion model consisting of a 3-layer MLP with 128 hidden units, preceded by a 2-layer
570 embedding module with 64 hidden units. This model is trained for 1 million gradient steps.

571 For both cases, we set the coefficient of the contrastive improvement loss $L_{\text{contrast}} = \max\{\bar{f}_0; L_{\text{prior}}$
572 $L_{\text{post}}\}g$ to 0.1. The key hyper-parameters are summarized in Table A10.

573 D.2.3 Policy

574 For the DP-based policy, we adopt the same architecture as the planner described earlier for Cheetah,
575 Maze2D, Kitchen, Ant, and Walker. For LIBERO, we use a Transformer-based noise predictor with a
576 decoder architecture comprising 12 layers, 12 attention heads, and a hidden embedding dimension of
577 768. Following DP [6], we apply dropout with a rate of 0.1 to both the input embeddings and attention
578 weights. The number of diffusion timesteps is 500. When conditioning is used, we incorporate a
579 Transformer encoder with 4 layers to encode the condition tokens, which include a sinusoidal timestep
580 embedding and projected observed trajectory tokens (all mapped to the same embedding dimension).
581 In this encoder-decoder setup, causal masking is applied to ensure autoregressive generation. In the
582 unconditioned case, we prepend the sinusoidal timestep embedding to the input sequence and use a
583 BERT-style encoder-only Transformer. All environments (Cheetah, Ant, Kitchen, Maze2D, Walker,
584 and LIBERO) are trained using the AdamW optimizer with a learning rate of 10^{-4} , weight decay
585 10^{-3} , $\beta_1 = 0.9$, and $\beta_2 = 0.95$. Layer normalization is applied before each Transformer block for
586 stability. The observation, planning, and action horizons follow the same settings used for the planner
587 in each environment.

588 For the IDQL-based policy, we align all hyperparameters for Cheetah and Ant with the original IDQL
589 implementation, using an observation, planning, and action horizon of 1. Hence, in IDQL-based ones,
590 we do not consider autoregressive modeling. Similarly, for both cases, we use consider the coefficient
591 before the contrastive improvement loss as 0.1.

592 D.3 Connection to Bayesian Filtering

593 In the absence of explicitly designed latent variables, our model can be interpreted as a form of
594 *Bayesian filtering* [70]. Under a general formulation of the hidden Markov model (HMM) [71]
595 with an additional latent dependency on observation ($C \setminus \mathbf{x}$), the latent process over C captures the
596 underlying stochasticity present in the demonstration data, which arises from both the environment
597 dynamics and the behavior policy. In this view, the latent variable acts as a compact and expressive
598 representation that summarizes the uncertainty in past observations, thereby improving the prediction
599 of future observations. This, in turn, facilitates more robust policy learning and planning in the
600 general settings.

601 E Extended Related Works

602 E.1 Diffusion Model-based Decision-making

603 Recent advances use diffusion models as the planner and policy for both reinforcement learning
604 (RL) and imitation learning (IL). RL agent aims to learn a policy that maximizes cumulative re-
605 wards through interaction with an environment [72]. The agent observes a sequence of transitions
606 $(\mathbf{s}_t; \mathbf{a}_t; r_t; \mathbf{s}_{t+1})$, where $\mathbf{s}_t \in S$ denotes the state, $\mathbf{a}_t \in A$ the action, $r_t \in \mathbb{R}$ the received reward,
607 and \mathbf{s}_{t+1} the next state. The goal is to learn a policy $\pi(\mathbf{a} | \mathbf{s})$ that maximizes the expected return:

608 $* = \arg \max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t r_t \right]$; where $\gamma \in [0, 1]$ is the discount factor. In contrast, IL [73] focuses
 609 on learning policies from expert demonstrations, often without access to the reward signal. A common
 610 approach is behavior cloning (BC) [74], which formulates IL as a supervised learning problem by
 611 maximizing the likelihood of expert actions given observed states, i.e., learning a policy $\pi_{\theta}(a|s)$ that
 612 closely imitates the expert policy $\pi_e(a|s)$.

613 **Diffusion Planner** Diffusion-based planning methods are commonly used to approximate the se-
 614 quence of future states and actions from a given current state. By leveraging the conditional generation
 615 capabilities of diffusion models—such as guidance techniques [29, 30]—these methods can generate
 616 plans (i.e., state trajectories) that satisfy desired properties, such as maximizing expected rewards.
 617 Taking Denoising Diffusion Probabilistic Models (DDPM [31])-based approaches as an example,
 618 these methods learn a generative model over expert trajectories $\mathcal{G} = \{(\mathbf{s}_0; \mathbf{a}_0); \dots; (\mathbf{s}_T; \mathbf{a}_T)\}$ by
 619 modeling a forward-noising process: $q(\mathbf{x}^k | \mathbf{x}^{k-1}) = \mathcal{N}(\mathbf{x}^k; \beta_k \mathbf{x}^{k-1}; (1 - \beta_k) \mathbf{I})$, and a parame-
 620 terized denoising model $p_{\theta}(\mathbf{x}^{k-1} | \mathbf{x}^k)$ to reverse the process. Here, k denotes the diffusion step, \mathbf{x}^0
 621 is a clean sub-sequence sampled from the expert trajectory \mathcal{G} , and β_k controls the variance schedule
 622 at step k .

623 During inference, trajectories are generated by starting from Gaussian noise and iteratively denoising
 624 through the learned reverse process. This generation can be optionally conditioned on the initial state
 625 or other guidance signals \mathbf{y} , such as rewards, goals, or other constraints: $\hat{\pi}_{\theta}(\cdot | \mathbf{s}_0; \mathbf{y})$.

626 These methods generally fall into two main categories: (1) learning a joint distribution over state-
 627 action trajectories, as in Diffuser [5], or (2) learning only state trajectories via diffusion and using an
 628 inverse dynamics model to recover actions, as in Decision Diffuser (DD) [55]. Beyond these, several
 629 variants extend diffusion-based planning in different directions. For example, Latent Diffuser [75]
 630 plans in a high-level latent skill space to improve generalization and LDP [54] plans with high-
 631 level latent actions directly from high-dimensional action-free demonstrations. Other approaches
 632 incorporate multi-task context to enhance adaptation and performance in unseen tasks, including
 633 MetaDiffuser [27], AdaptDiffuser [76], and MTDiff-p [77]. In addition, recent efforts have explored
 634 various extensions of diffusion planning, such as ensuring safety during generation [78], handling
 635 multi-agent scenarios [79, 80], learning skills [81], and application in RL from human feedback
 636 (RLHF) [82].

637 **Diffusion Policy** In contrast to diffusion-based planners, Diffusion Policy methods directly parame-
 638 terize the policy $\pi_{\theta}(a|s)$ using diffusion models. For example, Diffusion Policy [6] uses a diffusion
 639 model to generate actions with expressive, multimodal distributions. DPPO [7] extends this idea by
 640 modeling a two-layer MDP structure, where the inner MDP represents the denoising process and the
 641 outer MDP corresponds to the environment. This framework enables fine-tuning of diffusion-based
 642 policies in RL settings. Another line of work integrates diffusion models with model-free methods
 643 for offline RL by using diffusion models as to model the action distributions [32–35].

644 Recent explorations have also aimed to unify diffusion-based planning and policy learning within a sin-
 645 gle framework. For example, the Unified Video Action model (UVA) [83] and Unified World Models
 646 (UWM) [84] leverage diffusion models to jointly model planning and action generation, demonstrat-
 647 ing scalability on large-scale robotic tasks with pre-training. In a similar spirit, Ada-Diffuser
 648 provides a general framework that can be integrated into both diffusion planners and diffusion-based
 649 policies. However, Ada-Diffuser differs in its explicit modeling of latent factors that influence
 650 the data generation process. By incorporating latent identification directly into the diffusion process,
 651 Ada-Diffuser enables more structured, context-aware decision-making in partially observable
 652 and dynamically changing environments.

653 E.2 Latent Belief State Learning in POMDP

654 In partially observable Markov decision processes (POMDPs), single-step observations are typically
 655 insufficient for making optimal decisions. A common strategy to overcome this limitation involves
 656 encoding an agent’s history, encoding past observations and actions into a belief state that captures
 657 a distribution over latent environmental states. Although such belief representations can, in theory,
 658 support optimal policy derivation [1, 17, 85], their exact computation depends on full knowledge of the
 659 transition and observation models. This requirement quickly becomes intractable in high-dimensional
 660 settings.

661 To address this, recent work has focused on learning approximate belief representations directly from
 662 data. Notable approaches include those using recurrent neural networks [18] and variational inference
 663 methods [19, 25], which enable agents to encode temporal structure and uncertainty into compact
 664 latent embeddings. These representations are then used to inform downstream policy learning,
 665 optimizing for cumulative rewards.

666 This direction also aligns with developments in meta-reinforcement learning and non-stationarity,
 667 where belief states or Bayesian embeddings are used to capture hidden task contexts. Agents trained
 668 across a distribution of tasks can use these latent variables to infer new environments and adapt
 669 quickly [13, 20, 60, 61, 14]. For example, MetaDiffuser [27] incorporates task context as conditioning
 670 input to diffusion-based decision models.

671 Our approach diverges from these by offering theoretical guarantees on the identifiability of latent
 672 factors from minimal temporal observations. Rather than depending on diverse multi-environment
 673 data, we introduce a framework that captures the full data generation process in RL using diffusion
 674 models. In contrast to MetaDiffuser, which assumes static task-level context, our model treats the
 675 latent context as a dynamic, time-evolving process that governs both environment transitions and
 676 agent behavior, capturing the underlying temporal structure of RL trajectories more faithfully.

677 E.3 Autoregressive Diffusion Models

678 To model temporal consistency and dynamics in sequential data such as videos and audios, recent
 679 work has incorporated autoregressive structures into diffusion models. These approaches differ in how
 680 they condition on prior time steps during generation and can be categorized into two main categories.
 681 **(1) Conditioning on clean (denoised) inputs** ([86–88]). At each time step t , the denoising model
 682 is conditioned on the previously denoised outputs $\bar{\mathbf{x}}_{<t}^0 \mathcal{G}$: $p_\theta(\mathbf{x}_t^{k-1} \mid \mathbf{x}_t^k; \mathbf{x}_{<t}^0)$, where \mathbf{x}_t^k is the
 683 current noisy input, and $\mathbf{x}_{<t}^0$ denotes the clean (fully denoised) observations from earlier time steps.
 684 **(2) Conditioning on noisy inputs** ([21–24]). These methods instead condition on previous time
 685 steps at their corresponding noise levels. This setting can be further divided into two cases: (a)
 686 *fully noisy conditioning* [21]: the model conditions on all prior time steps at the same noise level
 687 k : $p_\theta(\mathbf{x}_{<t}^{k-1}; \mathbf{x}_t^{k-1} \mid \mathbf{x}_t^k; \mathbf{x}_{<t}^k)$. (b) *partially noisy conditioning*: each previous time step $i < t$ is
 688 conditioned at its own noise level k_i , which may vary over time: $p_\theta(\mathbf{x}_0^{k_0-1}; \mathbf{x}_1^{k_1-1}; \dots; \mathbf{x}_T^{k_T-1} \mid$
 689 $\mathbf{x}_0^{k_0}; \mathbf{x}_1^{k_1}; \dots; \mathbf{x}_T^{k_T})$. Specifically, Diffusion Forcing (DF) [22] proposes a general framework in which
 690 each time step \mathbf{x}_t assigns an independent noise level. In contrast, other works adopt time-dependent
 691 noise schedules that vary with the temporal index [23, 24, 28].

692 To model the causal generative process of RL trajectories, our approach also employs time-dependent
 693 noise scheduling to capture temporal dynamics. However, unlike prior work, we further integrate
 694 the identification of latent factors directly into the denoising process. This is achieved through a
 695 structured reinforcement step during training and a zig-zag inference procedure at test time, enabling
 696 our model to more faithfully recover the underlying causal structure in sequential decision-making.

697 E.4 Summary

698 To sum up, we compare our approach with representative diffusion- and meta-learning-based base-
 699 lines (Table A11). Diffuser, DP, IDQL, and DD do not model or infer latent contexts; DF adopts
 700 autoregressive denoising but still lacks context inference. Meta-Diffuser, LILAC, and DynaMITE
 701 learn latents via meta-learning but omit our minimal-sufficient block design and backward refinement.
 702 LDCQ and LDP model only high-level latent actions/skills without explicit context identification. In
 703 contrast, our method jointly models latent factors, employs full autoregressive denoising with zig-zag
 704 sampling, and introduces a backward refinement mechanism that enables identifiable latent contexts.

Method	Latent Factors	AR Denoising	Min. & Suff. Obs.
Ours	Yes (dyn., rew., act.)	Yes	Yes (refine, zig-zag)
Diffuser / DP / DD / IDQL	No	No	No
DF	No	Yes	No
Meta-Diffuser / LILAC / DynaMITE	Yes (dyn., rew. only)	No	No
LDCQ	Yes (hi-level act.)	No	No
LDP	Yes (hi-level act.)	No	No

Table A11: Comparison with representative baselines on whether they model latent contexts, use autoregressive (AR) denoising, and enforce minimal & sufficient observation blocks.

705 F Benchmark Settings and Illustrations

706 F.1 Latent Change Factors Design

707 We consider the latent change factors on dynamics and rewards. We consider the Half-Cheetah
708 and Ant environments from the OpenAI Gym suite, which are widely used MuJoCo locomotion
709 benchmarks [89] for evaluating continuous control algorithms. In Half-Cheetah, the agent is a
710 planar bipedal robot with a 17-dimensional state space and a 6-dimensional continuous action
711 space, where the goal is to move forward by applying torques to six actuated joints. In Ant, a
712 quadrupedal robot operates in a 3D space with a 111-dimensional state space and an 8-dimensional
713 action space, requiring more complex coordination across its four legs. In both environments, the
714 reward encourages forward velocity while penalizing excessive control inputs and, in the case of Ant,
715 also promotes stable contact with the ground. We consider variants of the Half-Cheetah environment
716 to study changes in dynamics, specifically **Cheetah-Wind-E** and **Cheetah-Wind-S**, which introduce
717 external wind forces applied to the agent. In **Cheetah-Wind-E**, an opposing wind force is applied at
718 the beginning of each episode and remains constant throughout, defined as $f_w = 10 + 5 \sin(0.8t)$,
719 where t is the episode index. For this case, since \mathbf{c} change over episode, we use data from several
720 consecutive episodes to estimate it. In **Cheetah-Wind-S**, the wind force varies at every time step
721 according to the same formula $f_w = 5 + 5 \sin(0.5t)$, with t now representing the time step in
722 each episode. We also consider variations in the reward function. In **Cheetah-Dir-E**, the reward
723 depends on a time-varying goal direction, requiring the agent to alternate between moving forward
724 and backward. Specifically, the reward at episode t is defined as

$$r_t = d_t \ v_t \ 0.1k\mathbf{a}_t k^2;$$

725 where v_t is the agent’s forward velocity, \mathbf{a}_t is the action vector (torques applied), and $d_t \in [-1; +1]$
726 indicates the target direction at time t . The direction signal d_t changes, giving a non-stationary reward
727 function that challenges the policy to adapt to shifting goals. Specifically, we consider

$$d_t = (5 \ \text{sigmoid}(\frac{t-200}{T}));$$

728 where $\text{sigmoid}(\cdot)$ denotes the sigmoid function, σ controls the sharpness of the transition, and T determines
729 the switching period. This formulation induces a smooth periodic change in the preferred direction of
730 movement, requiring the policy to adapt to gradually shifting objectives.

731 We also consider a directional reward variant for the **Ant** environment, denoted as **Ant-Dir-E**, where
732 the agent is required to alternate its movement direction over time. The reward function at time step t
733 is defined as

$$r_t = (2d_t - 1) \ v_t^x \ 0.1k\mathbf{a}_t k^2;$$

734 where v_t^x is the velocity of the agent’s torso along the x-axis (forward direction), \mathbf{a}_t is the 8-
735 dimensional action vector, and $d_t \in [0; 1]$ is a smooth directional signal. Similarly, we define d_t
736 as:

$$d_t = (5 \ \text{sigmoid}(\frac{t-200}{T}));$$

737 where $\text{sigmoid}(\cdot)$ denotes the sigmoid function. This formulation causes the preferred movement direction
738 to alternate approximately every 100 steps. Notably, for these settings with periodic changes (i.e.,
739 where latent factors do not evolve at every timestep), we estimate the latent variables periodically and
740 perform refinement in the causal diffusion model only when changes are detected. This follows the
741 same overall framework, but operates at a coarser temporal resolution aligned with the latent change
742 frequency.




Figure A3: **Illustrations of the Benchmarks.** From left to right: Half-Cheetah, Ant, Walker, Franka-Kitchen, Maze2D, and LIBERO.




Figure A4: **Illustrations of RoboMimic Benchmark.**

743 **F.2 Overview on Other Benchmarks**

744 Fig. A3-A4 give the illustrations on the used benchmarks. Specifically, other than Cheetah and Ant
 745 we introduced before, for others, we consider the basic settings in offline RL. Specifically,

746 **Maze2D.** Maze2D tasks focus on goal-directed navigation in a 2D plane, where the agent must
 747 traverse a maze-like environment to reach specified targets. These settings are designed to evaluate
 748 an agent’s ability to reason spatially and follow optimal trajectories based solely on positional and
 749 velocity observations.

750 **Franka-Kitchen.** The Franka-Kitchen environment [50] involves a robotic arm interacting with a
 751 series of articulated objects in a realistic kitchen setting. Tasks are composed of multiple stages, such
 752 as opening doors or toggling switches, and are intended to assess an agent’s capability in handling
 753 long-horizon, multi-step manipulation.

754 **Walker.** The Walker2D environment features a two-legged robot that must learn to walk and balance
 755 using continuous torque control. The agent’s objective is to maintain forward motion while remaining
 756 upright, which requires learning dynamic stability and coordination.

757 **LIBERO [53].** The Libero benchmark offers a diverse set of continual learning tasks focused on
 758 object manipulation and generalization:

- 759 • **LIBERO-Object:** The robot is required to manipulate a variety of novel objects through
 760 pick-and-place operations. Each task introduces previously unseen objects, encouraging the
 761 agent to incrementally build knowledge about object-specific properties and behaviors.
- 762 • **LIBERO-Goal:** All tasks share a common object set and spatial layout, but vary in goal
 763 specifications. This setup tests the agent’s ability to continually adapt to new task intents
 764 and motion targets without changes in the visual scene.
- 765 • **LIBERO-Spatial:** Tasks involve repositioning a bowl onto different plate locations. Al-
 766 though the objects remain fixed, the spatial configurations vary across tasks, requiring the
 767 robot to incrementally acquire relational spatial understanding.

768 **RoboMimic.** RoboMimic [52] provides a set of manipulation tasks based on human teleoperation
 769 demonstrations, varying in difficulty and required precision:

- 770 • **Lift:** The robot arm is tasked with lifting a small cube off the table. This task serves as a
 771 foundational manipulation scenario focused on grasping and vertical motion.

- 772 • **Can:** The robot must retrieve a cylindrical can from a cluttered bin and place it into a
773 designated smaller container. This task introduces greater complexity due to object shape
774 and the need for accurate placement.
- 775 • **Square:** A fine-grained insertion task where the robot picks up a square nut and places it
776 onto a vertical rod. This is the most challenging of the three, requiring precise alignment
777 and control for successful completion.

778 G Other Details on Ada-Diffuser

779 G.1 Latent Action Planner

780 For the latent action planner, we align our settings with those used in LDP [54], specifically focusing
781 on learning directly from image-based demonstrations. We first use a variational autoencoder (VAE)
782 to extract latent representations \mathbf{Z} from raw images via image encoders. An inverse dynamics model
783 is then trained to recover actions \mathbf{a}_t from pairs of latent states $(\mathbf{z}_t; \mathbf{z}_{t+1})$. A planner is subsequently
784 trained to forecast future latents \mathbf{Z} .

785 In our framework, we treat the latent factors \mathbf{C} as high-level latent actions that influence the evolution
786 of \mathbf{Z} . These latent factors are jointly used with \mathbf{Z} to perform both inverse dynamics modeling and
787 latent forecasting, enabling structured planning in the latent space.

788 We follow the experimental settings established in LDP [54]. Specifically, we use expert demonstra-
789 tions alongside action-free and suboptimal demonstrations. All hyperparameters and architectural
790 choices for the diffusion models are kept identical to those used in the original LDP implementation.
791 We also directly utilize the pre-trained image encoder provided by LDP. The only modification in our
792 framework is the introduction of an additional latent factor \mathbf{C} trained by our latent factor identification
793 stage, which is incorporated into the model to enhance latent action planning.

794 G.2 Noise Scheduling

795 In the autoregressive setting, we consider a monotonic increasing denoising schedule $f k_1; \dots; k_T g$.
796 In practice, we use a linear schedule where $k_i = \frac{i}{T} K$, with K denoting the maximum diffusion
797 step used in both training and sampling. We segment the sequence into temporal blocks of length
798 $T_x + 1$ ($T_x = T_o$ in all settings), and slide the time window forward by one step at a time. This
799 design ensures that the denoising steps progressively increase across the block, aligning the diffusion
800 process with the underlying temporal structure. Such a schedule encourages early steps to rely more
801 on strong priors and later steps to refine based on more contextual information. Additionally, for
802 better illustration, Fig. A2 provides a detailed illustration of the zig-zag sampling process within a
803 temporal block of 4 timesteps.

804 H Specific Design Choices for Baselines

805 For all baselines, unless otherwise specified, we use the same set of diffusion parameters detailed in
806 Appendix D.2.2–D.2.3. Below, we provide additional details on how specific methods are evaluated.
807 While their diffusion backbones remain consistent as in Appendix D.2.2–D.2.3, these methods include
808 custom design choices and method-specific hyperparameters that are evaluated accordingly.

809 H.1 Details on LILAC and DynaMITE

810 In these settings, we extend both LILAC and DynaMITE by incorporating a context encoder to
811 infer latent context variables \mathbf{c}_t , following their respective designs. Both methods learn belief
812 state embeddings from historical observations. For a fair comparison, we use the same latent
813 identification network architecture as in our framework, but modify the inputs according to each
814 method’s assumptions.

815 Specifically, LILAC and DynaMITE condition their inference networks solely on the historical
816 trajectory $\mathbf{x}_{1:t}$, without access to current and future information. Additionally, consistent with
817 their original implementations, we do not include a separate prior head on top of the GRU; both
818 methods share the encoder for posterior inference and prior prediction. And the primary difference

819 (in implementation) between these two methods lies in the temporal context used: LILAC maintains
 820 the full belief over the entire history, i.e., it conditions on $\mathbf{x}_{1:t}$ to infer \mathbf{c}_{t+1} , while DynaMITE uses
 821 only the most recent context, i.e., it infers \mathbf{c}_{t+1} based solely on \mathbf{x}_t .

822 All other hyperparameters are aligned with those used in our Stage 1 training. The estimated context
 823 variables are then provided as additional conditioning inputs to the diffusion-based models.

824 H.2 Details on Diffusion Forcing

825 For Diffusion Forcing, we adopt the same autoregressive noise schedule as in our method, which
 826 accounts for causal uncertainty, similarly to the formulation in Eq. D.1 of [22], to ensure a fair
 827 comparison. Additionally, we use the Monte Carlo Guidance (MCG) mechanism introduced in [22]
 828 for Maze2D, following the original setup. For all other environments, we use the same classifier
 829 guidance scheme as the other baselines to maintain consistency in evaluation.

830 I Ablation Analysis

831 I.1 Training/Inference Time Analysis

832 We conduct all experiments on 4× NVIDIA A100 or 8× RTX 4090 GPUs, depending on the model
 833 scale and environment requirements. The main computational overhead in our framework arises from
 834 two components: (i) the latent factor identification network, and (ii) the denoise-and-refine steps
 835 in the diffusion model. During sampling, the additional cost comes from zig-zag latent exploration
 836 and latent variable sampling. However, these steps do not substantially increase either training or
 837 inference time.

838 To quantify this, we report the training and inference speed of our method compared to the base
 839 models DD and DP across all environments (Table A12). Our framework introduces only a moderate
 840 computational overhead — typically 1.2–1.3× the runtime of vanilla diffusion backbones, correspond-
 841 ing to roughly 20–30 This cost can be further reduced through parallel latent sampling, lightweight
 842 context encoders, or refinement only at inference. Moreover, we additionally evaluate a Picard-
 843 accelerated variant (Table A13, [90]), where iterative refinement is parallelized by conditioning
 844 each denoising step on previously denoised nodes. With Picard iteration, inference time drops to
 845 0.7–0.8× of our default iterative sampler while maintaining comparable performance, demonstrating
 846 the potential for further acceleration.

Environment	Training Time (sec/epoch)		Inference Time (sec/rollout)	
	Ours vs DD	Ours vs DP	Ours vs DD	Ours vs DP
Cheetah	72.1 / 60.1 (1.20)	69.8 / 58.4 (1.20)	1.51 / 1.26 (1.20)	1.57 / 1.22 (1.29)
Ant	79.5 / 64.3 (1.24)	76.0 / 62.0 (1.23)	1.67 / 1.36 (1.23)	1.72 / 1.32 (1.30)
Walker	85.3 / 67.1 (1.27)	81.5 / 64.2 (1.27)	1.83 / 1.45 (1.26)	1.88 / 1.38 (1.36)
Maze2D	90.2 / 72.0 (1.25)	88.3 / 69.2 (1.28)	1.94 / 1.54 (1.26)	2.01 / 1.46 (1.38)
Libero	104.0 / 81.0 (1.28)	102.1 / 78.0 (1.31)	2.18 / 1.72 (1.27)	2.31 / 1.64 (1.41)
Kitchen	117.8 / 88.1 (1.34)	115.3 / 85.0 (1.36)	2.45 / 1.83 (1.34)	2.55 / 1.73 (1.47)

Table A12: Training and inference time comparison for Ada-Diffuser-planning and Ada-Diffuser-policy variants. We report absolute times (sec/epoch or sec/rollout) and relative overheads.

Environment	Ours (sec)	Ours+Picard (sec)
Cheetah	1.51	1.15
Ant	1.67	1.25
Walker	1.83	1.40
Maze2D	1.94	1.47
Libero	2.18	1.62
Kitchen	2.45	1.84

Table A13: Picard-accelerated inference.

847 I.2 Ablation Results

848 I.2.1 Full Results Supplement to Table 2

849 Table A14 presents the full ablation results across all environments, as a supplement to Table 2.
850 Overall, the results highlight the importance of the two key components in causal diffusion modeling:
851 latent identification and autoregressive diffusion, both of which are critical for performance.

852 I.2.2 Noise Schedule: Linear vs. Logistic vs. Sigmoid

853 We adopt a linear noise schedule by default since any monotonic, bounded schedule suffices to model
854 the data-generation process and linear is simple and stable in practice. To validate this choice, we
855 compare linear, logistic, and sigmoid schedules on three representative tasks. As shown in Table A15,
856 performance remains stable across schedules with no significant differences, supporting our default
857 choice.

Task	Schedule	Performance
Cheetah	linear	68.9
	logistic	63.6
	sigmoid	70.4
Maze2D	linear	161.4
	logistic	157.6
	sigmoid	168.5
Franka-Kitchen	linear	0.70
	logistic	0.72
	sigmoid	0.66

Table A15: Ablation on noise schedules. “Performance” is the task score (higher is better for Maze2D/Kitchen; lower magnitude negative is better for Cheetah as per the benchmark).

858 I.2.3 Effect of Temporal Block Length on Latent Identification

859 We further analyze the impact of temporal block length on latent identification. As shown in Fig. A5,
860 the results are consistent with findings reported in the main paper. When the number of observations
861 is insufficient (e.g., 4), identification performance degrades. Performance improves when the block
862 length is in a moderate range (5–20), indicating that sufficient temporal context is beneficial. However,
863 using overly long blocks (> 20) introduces redundancy and increases optimization difficulty, which
864 in turn harms performance.

865 **Clustering** We assess whether the learned latent space organizes states by the underlying context
866 on the Cheetah wind-change task, where the ground-truth latent evolves as $f_w(t) = 5 + 5 \sin(0.5t)$.
867 We sample 1000 time steps, discretize $f_w(t)$ into five equal-frequency bins to define target clusters,
868 embed the corresponding observations into the 20-dimensional learned latent representation, and run
869 k -means with $k = 5$. We compare our method with LILAC and DynaMIE, together with an ablation
870 that without refinement. Results are given in Fig. A6.


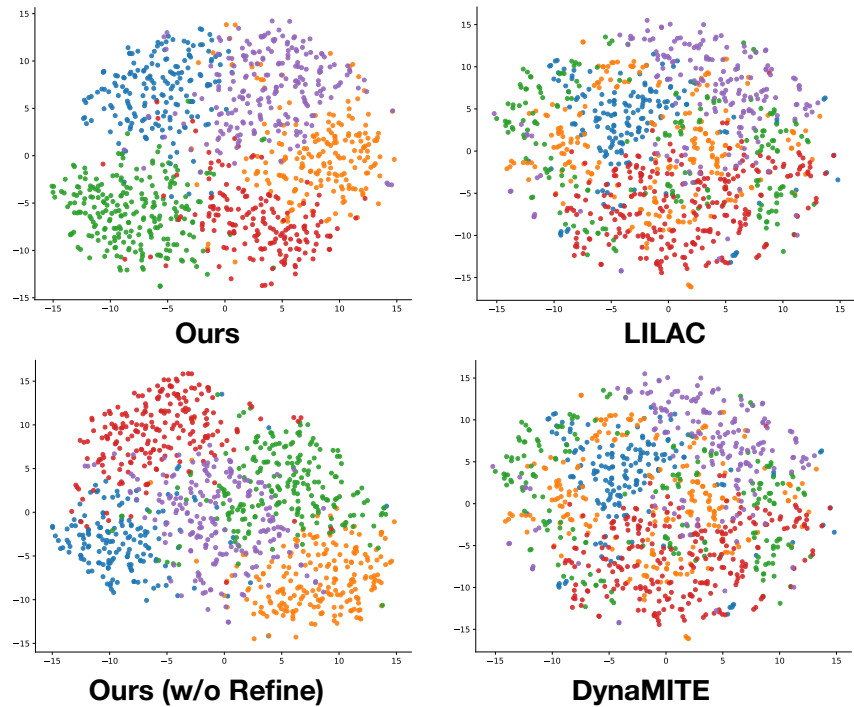


Figure A5: **Identification results (MSE of linear probing and R^2) versus the length of temporal blocks.** Left: Cheetah with time-varying wind; Right: Cheetah with time-varying rewards.



s

Figure A6: Clustering (t-SNE) results on Cheetah wind-change.

871 I.2.4 Latent Probing: Effect of Backward Refinement and Zig-Zag

872 To test whether backward refinement and zig-zag primarily help by correcting posterior mismatch,
 873 we perform a latent probing analysis on CHEETAH with changing wind. We linearly map the learned
 874 latent representation to the ground-truth wind variable using a simple least-squares probe (trained
 875 on a subset of blocks and evaluated on held-out blocks). Table A16 reports the mean squared error
 876 (MSE) of this probe for three variants: (i) the full model with backward refinement and zig-zag; (ii)
 877 without refinement; and (iii) without zig-zag.

Variant	MSE
Full (with refinement + zig-zag)	0.18
w/o refinement	0.28
w/o zig-zag	0.23

Table A16: Linear probing MSE for recovering the ground-truth wind latent on CHEETAH (changing wind). Lower is better.

878 **Analysis.** The full model achieves the lowest MSE, indicating more accurate recovery of the latent
879 wind. Removing backward refinement yields the largest degradation (0.18 / 0.28), consistent with
880 the role of refinement in letting future evidence within a block update the latent posterior and reduce
881 temporal lag. Disabling zig-zag also harms accuracy (0.18 / 0.23), suggesting that alternating
882 conditioning helps align the denoising trajectory with the latent dynamics rather than purely following
883 the forward temporal pass. Together, these results support our claim that both components reduce
884 posterior mismatch and improve latent identifiability, which in turn benefits planning and control in
885 settings with evolving hidden factors.

886 I.2.5 On the Effect of Planning and Execution Horizons: Long-horizon Planning

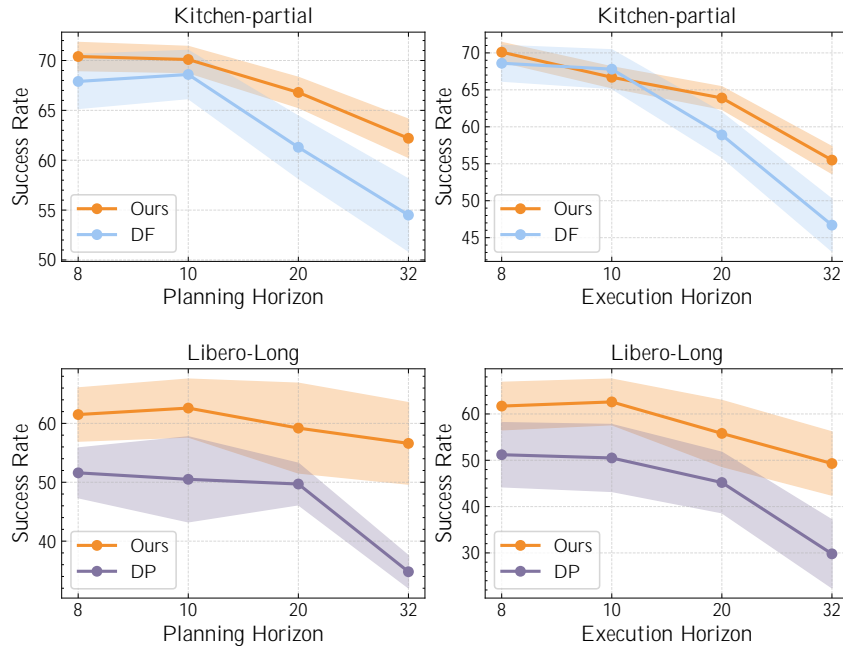


Figure A7: **Results with different planning and execution horizons.** We evaluate on Kitchen-partial and Libero-Long experiments.

887 We study the robustness of our approach under increased planning and execution horizons (T_p
888 and T_a). Specifically, we evaluate on two challenging tasks—Franka-Kitchen-Partial and Libero-
889 Long, where the original settings are Kitchen ($T_p = 16$, $T_a = 8$) and Libero ($T_p = 10$, $T_a = 8$).
890 Results are in Fig. A7. When we increase these horizons, we observe that the baselines, DP and
891 DF, suffer significant performance drops. In contrast, Ada-Diffuser maintains relatively high
892 performance. This demonstrates that modeling the underlying causal generative process, through
893 autoregressive structure and latent representations, enables better **long-horizon planning**. Although
894 we do not explicitly impose latent variables, our model implicitly learns representations that can track
895 stochasticity and support smooth control.

Environment	Diffuser	DF	MetaDiffuser	Diffuser + DynaMITE	Diffuser + LILAC	Ours	Ours + Meta	Oracle
Cheetah-Wind-E (c^s)	-120.4 ± 12.7	-105.8 ± 9.6	-89.7 ± 6.5	-79.2 ± 11.0	-95.3 ± 7.4	-68.9 ± 7.6	62.4 ± 3.9	57.8 ± 6.7
Cheetah-Wind-S (c^s)	-148.5 ± 9.8	-102.0 ± 10.2	-106.8 ± 11.4	-94.3 ± 9.6	-105.6 ± 14.5	-73.5 ± 8.7	-65.3 ± 11.2	-58.1 ± 9.0
Cheetah-Dir-E (c^r)	850.8 ± 54.2	902.1 ± 45.8	912.5 ± 37.9	930.4 ± 29.5	908.5 ± 37.6	943.3 ± 25.6	949.8 ± 24.1	962.1 ± 21.9
Cheetah-Vel-E (c^r)	-102.4 ± 18.2	-85.6 ± 18.3	-69.2 ± 7.5	-76.3 ± 11.7	-62.6 ± 11.1	-45.8 ± 9.5	-39.2 ± 7.6	-38.3 ± 8.9
Ant-Dir-E (c^r)	188.6 ± 39.2	195.4 ± 47.0	245.9 ± 41.0	262.8 ± 27.5	229.4 ± 32.6	285.3 ± 24.5	296.4 ± 22.2	300.7 ± 23.6

Table A4: **Results (average returns) on Ada-Diffuser-Planner with latent factors that affects dynamics and rewards.** c^s and c^r indicate the changes on dynamics and reward, E and S represent the episodic and time-step changes. The results are with 5 random seeds. The bold ones are the best-performing ones, excluding meta-learning and oracle ones.

Environment	DP	DP + DynaMITE	Ours + DP	Ours + DP (Oracle)	IDQL	IDQL + DynaMITE	Ours + IDQL	Ours + IDQL (Oracle)
Cheetah-Wind-E (c^s)	-104.8 ± 10.9	-72.2 ± 5.9	-58.5 ± 4.6	-52.0 ± 3.5	-97.5 ± 9.4	-59.0 ± 11.2	-48.5 ± 7.9	-41.6 ± 6.2
Cheetah-Wind-S (c^s)	-120.6 ± 11.5	-76.5 ± 15.6	-52.9 ± 9.8	-42.3 ± 6.7	-87.8 ± 12.2	-63.4 ± 6.7	-48.0 ± 7.2	-44.7 ± 6.1
Cheetah-Dir-E (c^r)	892.5 ± 60.8	949.6 ± 36.1	960.7 ± 40.2	972.4 ± 37.5	902.4 ± 45.2	938.6 ± 49.4	965.0 ± 37.5	969.8 ± 39.2
Cheetah-Vel-E (c^r)	-87.9 ± 6.5	-72.7 ± 5.8	-41.0 ± 7.2	-39.8 ± 6.7	-80.2 ± 11.4	-59.4 ± 6.5	-38.6 ± 7.7	-33.8 ± 6.5
Ant-Dir-E (c^r)	182.5 ± 41.2	275.2 ± 27.0	290.4 ± 49.4	312.5 ± 37.2	204.6 ± 25.6	269.3 ± 29.5	295.8 ± 32.7	309.6 ± 25.4

Table A5: **Results (average returns) on Ada-Di ffuser-Policy with latent factors.** c^s and c^r indicate the changes on dynamics and reward, E and S represent the episodic and time-step changes. The results are with 5 random seeds. The bold ones are the best-performing ones, excluding meta-learning and oracle ones.

Environment	Diffuser	DD	DF	LDCQ	Ours (DD)
Mixed	52.6 ± 2.3	75.2 ± 1.4	73.7 ± 1.9	73.3 ± 0.5	74.6 ± 1.6
Partial	55.8 ± 1.9	57.3 ± 1.2	68.6 ± 2.4	67.8 ± 0.8	70.1 ± 1.3

Table A6: Results (success rate (%)) on Ada-Diffuser-Planner without explicit latent factors on Franka-kitchen environment. The results are with 5 random seeds.

Environment	Diffuser	DD	DF	LDCQ	Ours (DD)
umaze	113.5 ± 2.8	114.8 ± 3.2	116.7 ± 2.0	134.2 ± 4.1	148.6 ± 3.7
medium	121.5 ± 5.6	129.6 ± 2.9	149.4 ± 7.5	125.3 ± 2.5	148.6 ± 3.1
large	123.0 ± 4.8	131.5 ± 4.2	159.0 ± 2.7	150.1 ± 2.9	161.4 ± 3.2

Table A7: Results on Ada-Diffuser-Planner without explicit latent factors on Maze-2D environment. The results are averaged across 5 random seeds.

Environment	Diffuser	DD	DF	LDCQ	Ours (DD)
medium-expert	106.2 ± 0.7	108.8 ± 2.0	105.4 ± 3.2	109.3 ± 0.4	115.7 ± 2.1
medium	79.6 ± 9.8	82.5 ± 1.6	66.2 ± 1.9	69.4 ± 2.4	83.6 ± 3.5
medium-replay	70.6 ± 0.6	75.0 ± 3.2	72.2 ± 2.6	68.5 ± 4.3	74.3 ± 2.8

Table A8: Results on Ada-Diffuser-Planner without explicit latent factors on Walker environment. The results are averaged across 5 random seeds.

Environment	DP	Ours (DP)
Spatial	78.3 ± 3.9	79.2 ± 4.2
Object	92.5 ± 2.6	93.4 ± 2.8
Long	50.5 ± 7.2	62.6 ± 4.9

Table A9: Results on Ada-Diffuser-Policy without explicit latent factors on Libero environment. The results are averaged across 5 random seeds.

Component	Type (i): Full Trajectory	Type (ii): State-Only
Model Backbone	1D U-Net [58]	Transformer (DiT)
Architecture	Kernel size: 5; channels: (1,2,2,2); base: 32	Hidden dim: 256; head dim: 32
# DiT Blocks	-	2 (Walker, Kitchen, Maze2D), 8 (LIBERO)
Optimizer	Adam, lr = 3×10^{-4}	Adam, lr = 3×10^{-4}
Batch Size	64	128
Training Steps	1M	1M
Diffusion Timesteps	150	200
Observation Horizon T_o	10	4 (Kitchen), 2 (LIBERO), 10 (others)
Planning Horizon T_p	16 (Cheetah), 32 (Ant)	16 (Kitchen), 10 (LIBERO), 32 (others)
Execution Horizon T_e	1	8 (Kitchen, LIBERO), 10 (others)
Guidance	CG, $\omega = 1.5$	CFG
Inverse Dynamics Model	-	2-layer embed (64), 3-layer MLP (128), 1M steps
Refinement Loss Coefficient	0.1	0.1

Table A10: Planner configurations for type (i): full trajectory generation and type (ii): state-only generation with inverse dynamics.

Cases	Cheetah-1	Cheetah-2	Ant	Maze2D	Walker	Kitchen	RoboMimic	LIBERO
Original	-73.5	-52.9	295.8	161.4	115.7	0.70	0.85	93.4
w/o refine	-82.0	-60.7	261.2	156.5	107.4	0.63	0.78	90.2
w/o zig-zag	-91.6	-56.1	258.3	147.6	107.9	0.59	0.75	91.6
same NS	-89.7	-62.4	259.7	140.1	105.8	0.56	0.72	85.2
random NS	-84.6	-62.9	266.4	146.3	109.1	0.61	0.76	88.5

Table A14: **Ablation on Design Choices.** We conduct ablation studies across a diverse set of tasks, including: Cheetah-Wind-S with a planner-based approach (denoted as Cheetah-1 in the table), Cheetah-Wind-S with a diffusion policy (Cheetah-2), Ant-Dir-E (policy, IDQL-based), Maze2D-Large (planner), Walker2D-Medium-Expert (planner), Kitchen-Partial (planner), LIBERO-Object (diffusion policy), and RoboMimic-Can.

896 **J LLM Usage Statement**

897 We disclose that LLMs were used solely to correct grammatical issues in this paper. It did not author
898 any sentence-level content. No part of the research ideas, experimental design, implementation, or
899 analysis relied on LLMs.