Hierarchical World Models with Implicit Dynamics

Gaoyue Zhou[1,2] **Yvonne Wu**[1] **Zichen Jeff Cui**[1]
Nicolas Ballas[2] **Mido Assran**[2] **Lerrel Pinto**[1] **Yann LeCun**[1,3]
[1]Courant Institute, New York University [2]Meta-FAIR [3]AMI Labs
Correspondence: gz2123@nyu.edu

Abstract

Planning with world models enables test-time adaptation for embodied agents, offering superior generalization over direct policy learning. However, long-horizon tasks remain hindered by cumulative errors and search complexity. Existing hierarchical methods often use rigid macro-actions or latent skills, leading to brittle representations and limited fine-grained control. We introduce **ID-WM (Implicit Dynamics World Models)**, a hierarchical world model that decouples state-space reachability from local control. Our framework utilizes a high-level model to sample feasible transitions across the state manifold and a low-level world model as an inverse dynamics solver to ground these transitions into precise actions. This allows the planner to identify reachable subgoals without being constrained by low-level execution details. By prioritizing local physical feasibility over global behavioral patterns, **ID-WM** synthesizes novel long-horizon plans from short segments, transcending the specific behavioral modes seen during training. We demonstrate its effectiveness across three navigation and manipulation suites, achieving a 38.5% improvement over state-of-the-art policy and diffusion-based planners. Crucially, by leveraging this hierarchical structure to mitigate compounding errors, our approach yields a 135% performance gain over single-level world model MPC, particularly in complex, unseen configurations for long-horizon tasks. Videos can be viewed at https://id-wm.github.io.

1 Introduction

World models (WM) [1, 2, 3] have emerged as a powerful paradigm for embodied decision-making [4, 5, 6, 7]. By learning predictive models of environment dynamics, they enable test-time reasoning and the synthesis of novel behaviors for unseen goals by searching in the action space [8, 9, 10, 11]. However, planning with WMs faces two fundamental challenges: 1) search complexity grows exponentially with planning horizon, and 2) prediction errors accumulate over long rollouts. Therefore, most WM-based planners are restricted to short-horizon tasks [8, 9], rely on simplified action abstractions [10], or operate in high-level spaces such as language, sacrificing fine-grained control [12].

Hierarchical learning offers a principled solution to this scalability challenge by decomposing long-horizon problems into manageable abstractions [13, 14, 15]. It has been extensively studied in Reinforcement Learning (RL) [16, 17, 18, 19], where hierarchy mitigates the credit assignment problem in long trajectories, and in imitation learning [20, 21], where it enables reusable behaviors. More recently, hierarchical structure has been introduced into diffusion-based planners to avoid directly modeling entire trajectories [22, 23, 24], improving scalability and generalization beyond fixed-horizon training data.

How can we enable the same capabilities for task-agnostic world models, enabling long-horizon zero-shot planning for unseen goals? A direct approach would be to simply increase the prediction horizon

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

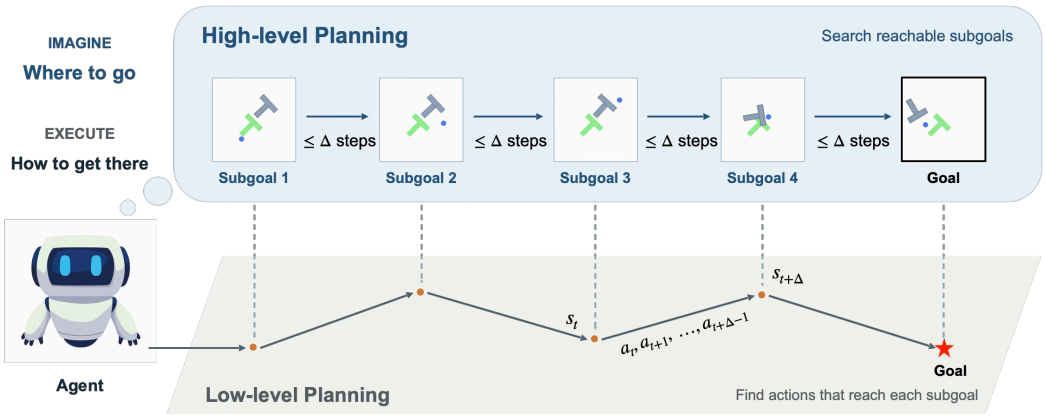


Figure 1: **Overview of ID-WM.** Our hierarchical planning framework decouples state-space reachability from local control, eliminating the reliance on task-specific value functions or pre-trained policies. Planning proceeds in two interleaved stages: (1) high-level planning, which discovers a sequence of temporally abstract subgoals spaced Δ steps apart, and (2) low-level control, which leverages a world model as an inverse dynamics solver to extract executable actions between each subgoal.

or action abstraction (e.g., frame-skipping) [25]. However, previous attempts at hierarchical world models typically depend on task-specific priors to function effectively, relying on pre-trained skills and low-level policies [19, 26] or rewards and value functions to verify subgoal feasibility [27, 28]. This dependency arises because planning is fundamentally adversarial: without these anchoring constraints, a powerful optimizer will inevitably exploit deficiencies in a learned world model, discovering "wormholes" or unrealistic transitions [29, 30]. To make hierarchical world models suitable for planning, we must construct a latent space that both condenses trajectory information to minimize search complexity, and remains inherently well-behaved to prevent the planner from exploiting unrealistic transitions.

To this end, we introduce **ID-WM**, a hierarchical world model with implicit dynamics grounded in the insight that effective planning requires decoupling topological reachability from execution details (Figure 1). **ID-WM** consists of a high-level diffusion-based world model that only implicitly models the dynamics by outputting states that are within a fixed distance away from the current state, functioning as a *neighboring state generator*, and a deterministic low-level model conditioned on both current state and action, acting as an *inverse dynamics solver*. Conditioned on the current state and a Gaussian noise, the high-level model leverages the generative power of diffusion to map a tractable unit Gaussian latent space to the manifold of feasible neighboring states. This formulation allows us to treat the noise as an implicit high-level action, enabling the planner to search over a well-behaved distribution instead of complex, discontinuous skill spaces to reason purely about "where to go," while the low-level model recovers the precise executable actions for "how to get there."

We summarize our contributions as follows:

- We introduce **ID-WM**, a hierarchical world model that decouples reachability from control, eliminating the need for task-specific value functions or pre-trained policies.
- We propose a corresponding hierarchical planning algorithm that bridges high-level abstraction with local dynamics, enabling robust, long-horizon decision-making.
- Across three navigation and manipulation environment suites, **ID-WM** shows 38.5% improvement for reaching arbitrary goals against powerful hierarchical offline reinforcement learning, imitation learning, and diffusion planning methods, and crucially, a 135% improvement over single-level world model and planning regimes.

2 Related Work

2.1 Hierarchical Learning

Hierarchical learning has long been a foundational pursuit in artificial intelligence, decomposing long-horizon problems into manageable abstractions [13, 14, 15]. This paradigm spans policy

architectures [31], reinforcement learning [16, 18, 27, 32], imitation learning [20], and planning [28, 33]. Existing work focuses on both *behavioral* (skills/options) [14, 19, 26, 34, 35, 36, 37] and *subgoal-centric* approaches [17, 28, 38]. However, these methods typically rely on coupled priors, such as rewards, learned value functions or pre-trained low-level policies [28, 32] to ensure feasibility, limiting generalization beyond the training distribution. In contrast, we remove this dependency, enabling robust hierarchical reasoning purely through implicit dynamics even when reliable value functions and policies are absent.

2.2 World Models for Embodied Planning

World models [1, 2, 3] have emerged as a foundational paradigm for embodied decision-making, enabling agents to reason about the future consequences of their actions. Recently, the field has shifted towards developing general-purpose world models [5, 8, 10, 39, 40, 41] capable of capturing broad environmental dynamics, diverging from traditional model-based reinforcement learning (MBRL) where learning task-specific reward or value functions is paramount [42, 43, 44, 45, 46, 47, 48]. Once these dynamics are learned, inference-time planning algorithms like the Cross-Entropy Method (CEM) [49] or gradient descent—are employed to derive behaviors zero-shot [8, 9, 11, 50]. However, the effectiveness of these approaches is often constrained to short horizons due to the computational complexity of search and the accumulation of prediction errors over time. While various strategies have been proposed to mitigate these issues, ranging from robust gradient-based optimization [30, 51] and efficient sampling routines [52] to temporal approximations for longer action sequences [10], they remain fundamentally flat in their temporal reasoning. Consequently, they struggle with long-horizon task solving, remaining susceptible to local minima and limited in their scalability to complex behaviors. Concurrently, HWM [25] builds a hierarchical latent world model by increasing the high-level frameskip and learning macro-actions; in contrast, **ID-WM** searches over the diffusion noise space, which is smooth and well-conditioned by construction.

2.3 Diffusion Models and Planning

Diffusion probabilistic models [53, 54, 55] excel at capturing complex, multimodal distributions via iterative denoising. Their efficacy is proven across high-dimensional modalities, including pixels [40, 41, 56, 57], language [58], and robot actions [59, 60]. Recently, diffusion models have been proposed as zero-shot planners for embodied tasks [61, 62, 63, 64]. These methods model entire trajectories directly, limiting generalization to out-of-distribution initial-goal pairs and constraining planning horizons to those seen during training. While hierarchical formulations mitigate these constraints, they rely on task-specific value functions [22, 23, 32], successful task trajectories [65], or assumed access to low-level primitives [24]. In contrast, we redefine diffusion as a *local generator* rather than a *global planner*. By delegating long-horizon reasoning to search, our method composes local jumps into novel trajectories without requiring task-specific priors.

3 Hierarchical World Models with Implicit Dynamics

Overview and Problem Formulation. We consider a goal-conditioned planning problem within a Markov Decision Process (MDP) defined by the tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P})$, where $\mathcal{S}$ is the continuous state space, $\mathcal{A}$ is the action space, and $\mathcal{P}(s_{t+1}|s_{\leq t}, a_{\leq t})$ represents the transition dynamics.

Following prior world modeling methods [8, 10, 11], we learn environment dynamics from a task-agnostic offline dataset without any reward labeling via a world model, and plan to reach arbitrary unseen goals at test time. More specifically, an agent is provided with an initial state o_0 and a goal state o_g at test time. The objective is to execute a sequence of actions $(a_0, a_1, \ldots, a_{T-1})$ such that the distance between the agent's final state o_T and the goal o_g is minimized. Unlike prior work, we do not assume a fixed task horizon T [8], focusing instead on enabling the agent to perform autonomous long-horizon reasoning and planning.

3.1 Hierarchical World Models with Implicit Dynamics

The core component of **ID-WM** is a hierarchical world model comprising a high-level diffusion model and a low-level inverse dynamics model. As illustrated in Figure 2, the high-level model operates as a *reachable state generator*, predicting a distribution of feasible future states conditioned

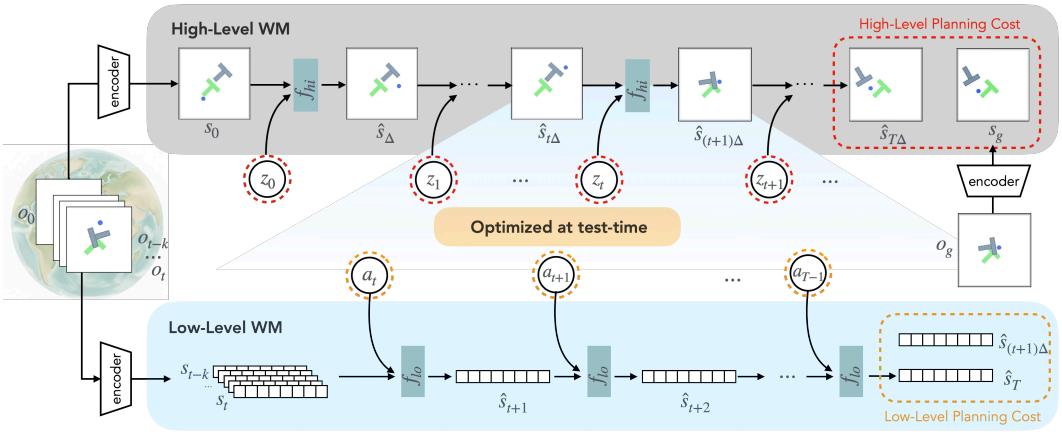


Figure 2: Architecture of **ID-WM**. Given an initial observation o_0 and a goal o_g, we optimize for both the sequence of subgoals $s_0, \ldots, s_{T\Delta}$ and the executable actions a_t in an interleaved fashion. Here, f_{hi} denotes the diffusion-based high-level world model and f_{lo} denotes the deterministic, transformer-based low-level world model, both of which predict in latent space.

on the current state. The low-level model acts as an *inverse dynamics solver*, grounding these predicted transitions into executable actions.

$$\begin{aligned} \text{High-level WM:} \quad & f_{hi}(s_{t+\Delta} \mid s_t, \Delta) \\ \text{Low-level WM:} \quad & f_{lo}(s_{t+1} \mid s_{t-k:t}, a_{t-k:t}) \end{aligned}$$

Specifically, the high-level world model conditions on the current state s_t and a temporal radius Δ, generating future states that fall within a Δ-step horizon. The low-level world model adopts the formulation from Zhou et al. [8], a transformer-based world model conditioning on a historical context of k past states and actions to infer immediate next steps.

Both levels of the hierarchy share a unified structure, consisting of the following components:

$$\begin{aligned} \text{Observation model:} \quad & s_t \sim \text{enc}_\theta(s_t \mid o_t) \\ \text{Transition model:} \quad & s_{t+1} \sim p_\theta(s_{t+1} \mid s_{\leq t}, c) \\ \text{Decoder model:} \quad & \hat{o}_t \sim q_\theta(o_t \mid s_t) \end{aligned}$$

where c is the context that the transition model takes in to make predictions. For f_{hi}, c is the temporal distance Δ, instructing how far away the subgoals should be. For f_{lo}, c is the sequence of executable actions the agent has taken up to time t (i.e., $a_{\leq t}$).

3.1.1 High-level Implicit Dynamics Model

To model the manifold of reachable states without restricting the agent to a fixed set of discrete skills, we formulate the high-level world model as a conditional diffusion probabilistic model (DDPM) [54]. Let x_0 represent the target future state $s_{t+\Delta}$ (at a fixed temporal offset Δ). The forward diffusion process $q(x_j|x_{j-1})$ gradually adds Gaussian noise to the data according to a variance schedule $\beta_1, \ldots, \beta_K$:

$$q(x_j|x_{j-1}) = \mathcal{N}(x_j; \sqrt{1-\beta_j}x_{j-1}, \beta_j \mathbf{I}) \tag{1}$$

This allows us to sample x_j at any arbitrary diffusion step j in closed form. We learn a conditional denoising network $\epsilon_\theta(x_j, j, s_t)$ to reverse this process, where s_t is the current state to condition the generation. The training objective is to predict the noise ϵ added to the clean data x_0:

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{x_0, \epsilon \sim \mathcal{N}(0,I), j}\left[\|\epsilon - \epsilon_\theta(x_j, j, s_t)\|^2\right] \tag{2}$$

Sampling as Implicit Dynamics (DDIM). To represent the high-level decision at environment time t, we let $z_t \in \mathbb{R}^d$ denote the initial noise vector x_K used for sampling the next subgoal. During inference, we employ Denoising Diffusion Implicit Models (DDIM) [55] to deterministically map

this noise vector $z_t \sim \mathcal{N}(0, \mathbf{I})$ to a generated subgoal $\hat{s}_{t+\Delta}$. The deterministic reverse step from diffusion step j to $j-1$ is given by:

$$x_{j-1} = \sqrt{\bar{\alpha}_{j-1}} \underbrace{\left(\frac{x_j - \sqrt{1-\bar{\alpha}_j}\epsilon_\theta(x_j, j, s_t)}{\sqrt{\bar{\alpha}_j}} \right)}_{\text{predicted } x_0} + \sqrt{1-\bar{\alpha}_{j-1}}\epsilon_\theta(x_j, j, s_t) \tag{3}$$

where $\alpha_j = 1 - \beta_j$ and $\bar{\alpha}_j = \prod_{i=1}^{j} \alpha_i$. This formulation effectively parameterizes the high-level transition dynamics as a deterministic function $\mathcal{G}_\theta(s_t, z_t)$.

Implicit High-Level Action Space. A key insight of our approach is to treat the initial noise vector z_t as an *implicit high-level action*. Unlike discrete skill indices or learned latent codes which may suffer from mode collapse or topological discontinuities, illustrated in Section 6.16, the noise space $\mathcal{Z} = \mathbb{R}^d$ is inherently smooth and well-behaved (isotropic Gaussian). By optimizing over z_t, the planner searches over the standard Gaussian latent space, where any sample z maps to a physically valid reachable state on the data manifold via the deterministic DDIM mapping.

3.1.2 Low-level World Model

The high-level World Model (WM) provides long-horizon guidance by identifying target states, yet it lacks the granular transition dynamics necessary to prescribe executable actions. To this end, we employ a transformer-based deterministic low-level WM following Zhou et al. [8], which functions as an inverse dynamics model through planning.

Given a history of states $s_{t-k:t}$ and actions $a_{t-k:t}$, the low-level WM predicts the next state $\hat{s}_{t+1} = f_{lo}(s_{t-k:t}, a_{t-k:t})$. The model is trained via supervised learning to minimize the prediction error:

$$\mathcal{L} = \|\hat{s}_{t+1} - s_{t+1}\|_2 \tag{4}$$

Consistent with Zhou et al. [8], we utilize an MLP action encoder to process raw action sequences before they are input into the transformer backbone.

At inference, we leverage this model for inverse dynamics through planning. Given a target state s^*, we search for an optimal action sequence $(a_0, \ldots, a_{H_{lo}-1})$ over a low-level planning horizon H_{lo} such that the planning cost $\|\hat{s}_{H_{lo}} - s^*\|_2$ is minimized.

3.2 Hierarchical Planning with ID-WM

At inference, given an initial state o_0 and a goal state o_g, we employ hierarchical planning to determine a sequence of actions that reaches o_g. We formulate this as a Model Predictive Control (MPC) problem, where each iteration consists of two phases: subgoal generation and subgoal realization. The overall hierarchical planning procedure is detailed in Algorithm 1. We also include a full version in Algorithm 2. Throughout, Δ denotes the temporal radius of each subgoal (the maximum number of environment steps between consecutive subgoals), H_{hi} the number of subgoals per high-level plan, and H_{lo} the low-level planning horizon used to reach each subgoal.

Subgoal Search via High-level WM. In the first phase, we identify a sequence of latent subgoals $s_0^*, \ldots, s_{H_{hi}}^*$ such that $s_0^* = s_0$ and $s_{H_{hi}}^* = s_g$. To ensure that each subgoal remains within the reachable manifold of the low-level WM given its search budget, we constrain the distance between consecutive subgoals s_i^*, s_{i+1}^* to be no more than Δ environment steps. We optimize the latent noise variables $z_0, \ldots, z_{H_{hi}-1}$ using the Cross-Entropy Method (CEM) to minimize the following planning cost:

$$\mathcal{L} = \|\hat{s}_{H_{hi}} - s_g\|_2 + \alpha \mathcal{L}_{reg}(Z) \tag{5}$$

where $\mathcal{L}_{reg}$ is a regularization term designed to keep the sampled noise within the support of the standard Gaussian distribution. This prevents the planner from exploiting adversarial regions of the latent space that would result in out-of-distribution transitions during DDIM sampling.

Subgoal Realization via Low-level WM. We similarly employ the CEM optimizer to identify the executable action sequences required to reach each intermediate subgoal. The agent plans and executes actions to reach the immediate subgoal, proceeding sequentially through the high-level

trajectory. Upon completing the sequence for all subgoals generated by the high-level WM, the agent initiates a new MPC iteration to replan the high-level trajectory from its current state.

Latent Noise Regularization. While the latent space is theoretically structured as a unit Gaussian, unconstrained planning can favor adversarial latents that produce unrealistic state transitions. To mitigate this, we incorporate a regularization term in the high-level planning cost:

$$\mathcal{L}_{reg} = \mathbb{E}_i \left[(\max(0, |z_i| - \tau))^p \right] \tag{6}$$

This term penalizes latent vectors whose magnitude exceeds a threshold τ, keeping the optimization within high-density regions of the prior. We examine the effect of noise regularization in Section 6.9.

Algorithm 1 Hierarchical Planning with **ID-WM**

Require: Initial state s_0, goal state s_g, high/low models f_{hi}, f_{lo}, planning costs C_{hi}, C_{lo}, planning horizons H_{hi}, H_{lo}.

$s_t \leftarrow s_0$
while goal s_g not reached **do**
  `# High-Level Plan`
  $(\hat{s}_1, \ldots, \hat{s}_{H_{hi}}) \leftarrow$ Planner$(s_t, s_g, f_{hi}, C_{hi})$
  `# Low-Level Plan and Execution`
  for $i = 1$ to H_{hi} **do**
    $(a_1, \ldots, a_{H_{lo}}) \leftarrow$ Planner$(s_t, \hat{s}_i, f_{lo}, C_{lo})$
    Execute actions $(a_1, \ldots, a_{H_{lo}})$ in the environment
    $s_t \leftarrow$ observe new state
  end for
end while

4 Experiments

Our experiments are designed to answer the following key questions:

- **Subgoal Generation:** Can **ID-WM** reliably generate diverse and feasible reachable states?
- **Generalization:** Can **ID-WM** solve long-horizon planning tasks to reach arbitrary, zero-shot goals at test time?
- **Scaling:** How does the performance of **ID-WM** scale across environments of increasing complexity?
- **Visual Planning:** Is **ID-WM** effective for long-horizon planning directly from high-dimensional visual observations?

4.1 Environments

We evaluate **ID-WM** on three navigation and manipulation environment suites as shown in Figure 3. The environments are Wall, PushT, and Maze, with task horizon ranging from hundreds to thousands of environment steps. A full description of the environments is detailed in Section 6.4.

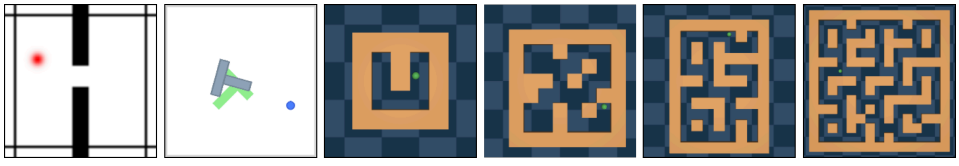


Figure 3: We evaluate **ID-WM** on three environment suites (Wall, PushT, Maze); from left to right: Wall, PushT, and the four maze sizes (U-, Medium, Large, Giant).

4.2 Baselines

We compare **ID-WM** with five state-of-the-art offline learning methods including imitation learning (GCBC), offline RL (HIQL), diffusion planners (Diffuser, Hierarchical Diffuser), and planning with world models (Single-level WM with MPC). See Section 6.7 for full details.

4.3 Generating Neighboring States with ID-WM

We first investigate the extent to which each level of **ID-WM** accurately captures the underlying data distribution and predicts future states across varying temporal granularities. For the high-level WM, we evaluate whether sampling from the latent prior yields diverse and reachable future states.

In Figure 4, we sample 100 latent noise vectors $z \sim \mathcal{N}(0, I)$ and visualize the resulting generated states $\hat{s}_\Delta = f_{hi}(s_0, z, \Delta)$ reachable within a Δ-step horizon. For each Δ, we overlay all 100 generated samples to visualize the predicted state distribution, explicitly marking the agent's position with a black dot for clarity.

As illustrated in Figure 4, the high-level WM generates diverse and physically reachable states for both the Wall and PushT environments. In the Wall environment, the dispersion of agent positions increases with Δ, reflecting a larger reachable radius as the temporal budget grows. Similarly, in the PushT environment, we observe an increasingly diffused distribution of both the agent and the T-block. The smooth coverage of object configurations demonstrates that the model successfully captures the complex, multi-modal interactions between the agent and the object. This high-fidelity distribution modeling serves as a necessary prerequisite for planning, enabling the system to generate diverse trajectories and generalize effectively to unseen goals in downstream tasks.

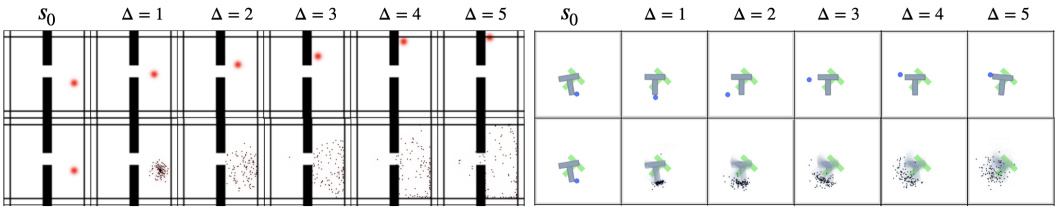


Figure 4: Diverse neighboring states generated with the high-level WM of **ID-WM**.

4.4 Hierarchical Planning with ID-WM

Can **ID-WM** solve long-horizon planning tasks to reach arbitrary, zero-shot goals at test time? To this end, we conduct hierarchical planning on three environment suites.

For all environments and baselines, we randomly sample 50 pairs of initial and goal states, with sampling method detailed in Section 6.5, and report the average performance across these instances. For the PushT task, we report two metrics: the success rate and the final coverage of the T-block relative to the target configuration. A trial is considered successful only if both the agent and the T-block reach their respective goal positions. In contrast, the final coverage metric provides a more granular measure of performance by focusing specifically on the object's alignment. To further evaluate generalization capabilities, we test PushT under two settings: a standard regime and a data-augmented regime featuring random spatial transformations and rotations. All experiments are conducted with a maximum horizon of 2000 environment steps. For all experiments, **ID-WM** uses $\Delta = 5$ (25 environment steps with frameskip) and $H_{hi} = 5$ subgoals per high-level plan, without per-task tuning; the planner replans after each executed plan (up to 15 replans), so no prior knowledge of the goal distance is required. We ablate Δ and the number of subgoals in Section 6.10.

In Table 1, **ID-WM** outperforms all baselines, achieving a 38.5% overall improvement and 135% gain over single-level WM with MPC (SL-WM). The gap is most evident in PushT with random goals, where **ID-WM** is the only approach exceeding a 30% success rate without augmentation. While SL-WM performs well on shorter horizons (Maze, Wall), it fails on PushT's long-horizon requirements (often > 100 steps) due to exponential search complexity and compounding prediction errors. Unlike the 25-step benchmark in Zhou et al. [8], our results show that standard WM planning cannot solve these extended tasks, a limitation our hierarchical approach addresses. We visualize the high-level plans in Figure 5 and examine subgoal quality and diversity in Sections 4.3 and 6.15.

Furthermore, the remaining baselines exhibit distinct failure modes. GCBC performs well on in-distribution goals (e.g., fixed-goal PushT) and in environments with dense data coverage (Medium Maze, Wall) but degrades significantly on out-of-distribution (OOD) goals. HIQL shows potential for OOD generalization but demands dense data coverage to accurately learn value functions, reflected in its performance gap between PushT and PushT-Aug. Diffuser-based methods struggle with dataset

Table 1: Performance comparison across environments. We report success rate (SR) and, for PushT, final coverage (Cov), as mean ± std over three seeds. PushT is evaluated with a fixed goal (Fix) and random goals (Rand.). Best results per row are in bold.

Environment	Metric	HIQL	GCBC	Diffuser	H-Diffuser	SL-WM	Ours
Medium Maze	SR	0.83 ±0.01	0.80 ±0.00	0.80 ±0.13	0.92 ±0.03	0.97 ±0.01	**0.99 ±0.01**
Wall	SR	0.62 ±0.12	**1.00 ±0.00**	0.07 ±0.08	0.10 ±0.02	0.88 ±0.02	**1.00 ±0.00**
PushT (Fix)	SR	0.45 ±0.07	**1.00 ±0.00**	0.00 ±0.00	0.01 ±0.01	0.19 ±0.02	**1.00 ±0.00**
	Cov	0.50 ±0.05	**0.85 ±0.04**	0.10 ±0.01	0.35 ±0.09	0.34 ±0.03	**0.85 ±0.01**
PushT (Rand.)	SR	0.02 ±0.02	0.21 ±0.01	0.00 ±0.00	0.01 ±0.01	0.21 ±0.04	**0.77 ±0.01**
	Cov	0.22 ±0.03	0.31 ±0.01	0.05 ±0.01	0.20 ±0.08	0.27 ±0.02	**0.58 ±0.01**
PushT-Aug (Fix)	SR	0.87 ±0.01	**1.00 ±0.00**	0.01 ±0.02	0.55 ±0.14	0.10 ±0.04	**1.00 ±0.00**
	Cov	0.81 ±0.05	0.83 ±0.01	0.31 ±0.03	0.74 ±0.16	0.43 ±0.01	**0.84 ±0.01**
PushT-Aug (Rand.)	SR	0.32 ±0.06	0.15 ±0.01	0.01 ±0.01	0.27 ±0.03	0.10 ±0.04	**1.00 ±0.01**
	Cov	0.59 ±0.03	0.25 ±0.01	0.14 ±0.01	0.72 ±0.19	0.17 ±0.01	**0.77 ±0.01**

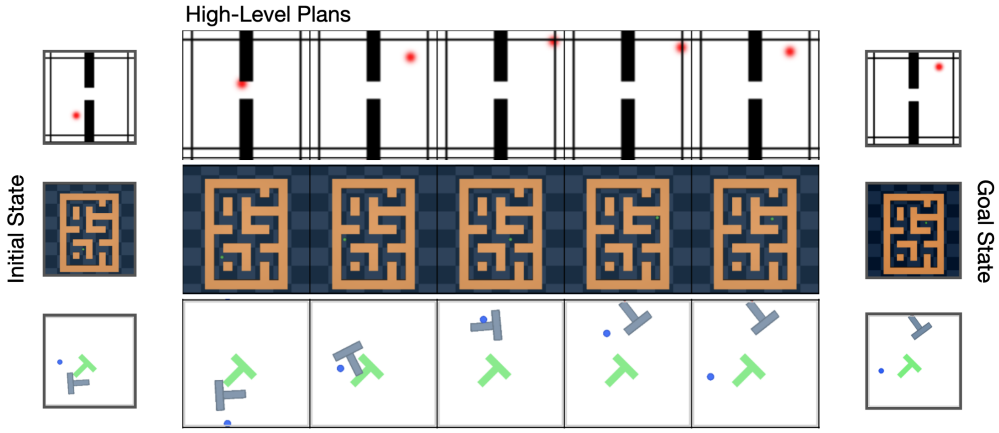


Figure 5: High-level subgoals planned by **ID-WM**. The planned sequence successfully guides the agent to the target goal while ensuring that consecutive subgoals satisfy local reachability constraints.

noise; despite adding replanning (Section 6.14.8), performance remains poor. While H-Diffuser offers marginal improvements, it still relies heavily on dense data. Even with data augmentation (Section 6.6), H-Diffuser's success rate remains low, primarily because it struggles to jointly attend to both the object and the agent's final position.

Comparison with hierarchical latent world models. We additionally compare with HWM [25], a hierarchical world model that increases the frameskip of its high-level model and compresses action sequences into macro-actions with a transformer-based action encoder. We compare with the DINO-WM variant of HWM. As shown in Table 2, on PushT, which requires long-horizon object manipulation, **ID-WM** outperforms HWM by a large margin on all variants. This highlights the benefit of disentangling state reachability from fine-grained control using diffusion noise as an implicit latent action.

Table 2: Success rate of **ID-WM** and HWM on PushT environments with fixed (Fix) and random (Rand.) goals.

	PushT		PushT-Aug	
Method	Fix	Rand.	Fix	Rand.
HWM	0.38	0.20	0.42	0.16
ID-WM	**1.00**	**0.77**	**1.00**	**1.00**

Unlike approaches that explicitly model dataset behavior, **ID-WM** is not constrained by the underlying training distribution. By leveraging full planning over the environment dynamics, our method

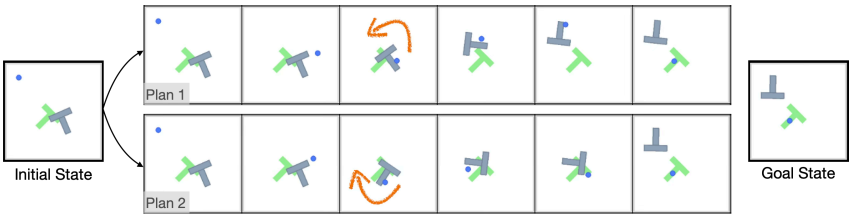


Figure 6: Diverse plans generated by **ID-WM**. The top and bottom rows illustrate distinct strategies: counter-clockwise and clockwise rotations of the T-block, both resulting in successful task completion.

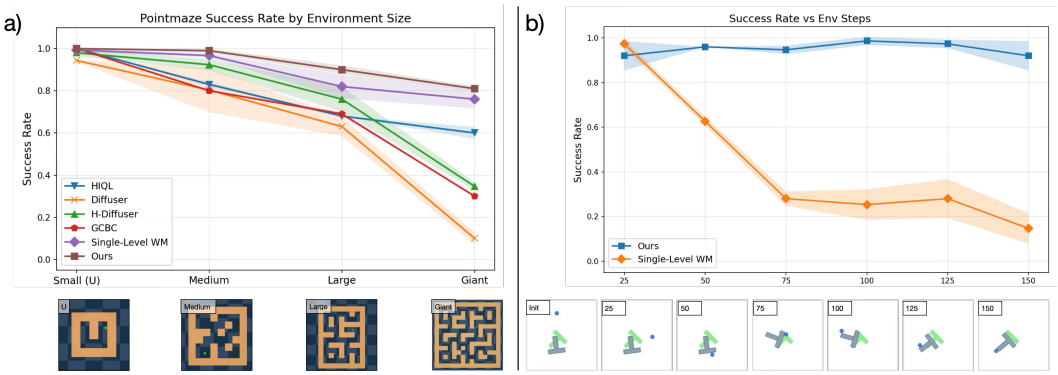


Figure 7: Planning performance with increasing task complexity. In a), we test **ID-WM**'s ability on Maze with increasing sizes. In b), we test **ID-WM** on increasingly longer horizon PushT tasks.

is inherently capable of finding diverse solutions for a given goal. Figure 6 demonstrates this by showing two different subgoal sequences generated by **ID-WM**. For instance, Plan 1 (top row) rotates the T-block counterclockwise, while Plan 2 (bottom row) rotates it clockwise. Ultimately, both planned trajectories successfully achieve the final objective via our hierarchical framework.

4.5 How does ID-WM scale with task complexity?

We investigate whether **ID-WM** scales to tasks of increasing complexity across two dimensions: environment size and goal horizon.

Scaling with Maze Size. We evaluate **ID-WM** across four mazes of increasing difficulty: U-Maze (3×3), Medium Maze (6×6), Large Maze (7×10), and Giant Maze (13×13) (Figure 7a). Each dataset contains 2,000 random 100-step trajectories. Given a fixed dataset size, data coverage naturally becomes sparser as maze size increases.

While all methods achieve near-perfect success on the U-Maze, baselines degrade significantly in larger environments (Figure 7a). In contrast, **ID-WM** maintains the highest performance on the Large Maze. We hypothesize this is because **ID-WM** combines local dynamics modeling with high-level planning, effectively "stitching" together regions of sparse data. Unlike GCBC or Diffuser (which model entire trajectories simultaneously) or HIQL (which requires dense data for global value functions), **ID-WM** does not rely on global trajectory coverage. While H-Diffuser uses abstraction, it lacks an explicit dynamics model, resulting in brittle plans that require frequent replanning. By grounding its hierarchy in explicit local dynamics, **ID-WM** generates robust, scalable plans.

Scaling with Goal Horizon. Building on the PushT planning benchmark introduced by Zhou et al. [8] (which evaluates goals 25 steps away), we increase complexity by sampling goals from the validation set that are $d \in \{25, 50, \ldots, 150\}$ steps away.

We compare **ID-WM** against a single-level world model baseline that performs full-horizon CEM search and replanning at every step. While the total number of executed actions is identical, **ID-WM** achieves a near-perfect success rate across all horizons as shown in Figure 7b). This confirms our main PushT results (Table 1), demonstrating that **ID-WM** maintains precise goal-reaching capabilities even when goals are over 150 steps away.

4.6 Vision-based Planning with ID-WM

To evaluate **ID-WM** in high-dimensional observation spaces, we test it on a vision-based PushT task. We use the state-based PushT dataset augmented with four 90° rotations for spatial robustness. Observations ($3 \times 64 \times 64$ images) are encoded into $4 \times 8 \times 8$ latents using a pretrained Stable Diffusion VAE [66]. We compare against AVDC [64], a language-conditioned video planner that generates all subgoals simultaneously (details in Section 6.7).

As shown in Figure 8, **ID-WM** generates reliable, sequential subgoals that ensure steady progress. In contrast, AVDC fails to produce meaningful plans, likely because the initial-goal pairs are out-of-distribution. This shows the generalization advantage of iterative subgoal search over simultaneous generation—which relies on trajectory-level dataset coverage—even in compressed visual spaces.



Figure 8: Subgoals generated by **ID-WM** on a visual PushT task. The first column denotes the starting state, and the last shaded column denotes the desired goal state. **ID-WM** generates reliable, sequential subgoals that drive steady progress toward the goal state without any degenerate states.

Quantitative evaluation. To complement the qualitative comparison, we report success rates on vision-based PushT in Table 3. We compare against two additional baselines: DINO-WM [8], a flat single-level world model planner, and HWM [25], a concurrent hierarchical latent world model that uses a larger frameskip for its high-level model. In this setting, **ID-WM** uses DINO-WM as its low-level world model, so the comparison between DINO-WM and **ID-WM** isolates the gain provided by the hierarchy. We evaluate four settings: for $d \in \{25, 50, 75\}$, the goal is sampled from a held-out validation set and is known to be d steps away; for Random Goal, both the initial and goal states are sampled randomly from the environment. **ID-WM** achieves the best performance on the most challenging long-horizon settings ($d = 50$, $d = 75$, and Random Goal), demonstrating its effectiveness for long-horizon vision-based planning.

Table 3: Success rate on vision-based PushT. For $d \in \{25, 50, 75\}$ the goal is d steps away; Random Goal samples both the initial and goal states randomly.

Method	$d = 25$	$d = 50$	$d = 75$	**Random Goal**
DINO-WM	0.84	0.16	0.00	0.00
HWM	**0.96**	0.46	0.14	0.02
ID-WM	0.90	**0.60**	**0.50**	**0.26**

5 Conclusion

We present **ID-WM**, a hierarchical world modeling and planning framework that decouples temporally extended subgoal generation from local reachability. By treating diffusion noise as implicit actions, we constrain the planning search space to a unit Gaussian, ensuring reliable subgoal generation. **ID-WM** significantly outperforms state-of-the-art baselines and demonstrates high-quality vision-based subgoal generation. Ultimately, grounding hierarchical abstraction in explicit local dynamics represents a promising step toward robust, long-horizon decision-making.

Limitations. First, **ID-WM**'s reliance on search-based optimization within a diffusion latent space incurs higher inference times than reactive policies or deterministic world models. Future work will explore amortized or gradient-based planning to mitigate this bottleneck. Second, while our hierarchy prevents planning complexity from scaling with trajectory length, it still scales with state-space dimensionality. Integrating compact, structured representation learning could further streamline this planning process for highly complex environments.

Acknowledgments and Disclosure of Funding

We would like to thank Amir Bar, Alfredo Canziani, Yilun Kuang, Hengkai Pan, Mike Rabbat, Ying Wang, and Wancong Zhang for their valuable discussion and feedback. This work was supported by AFOSR under grant FA95502310139, the Institute of Information & Communications Technology Planning & Evaluation (IITP) with a grant funded by the Ministry of Science and ICT (MSIT) of the Republic of Korea in connection with the Global AI Frontier Lab International Collaborative Research Agreement (No. RS-2024-00469482 & RS-2024-00509279).

References

[1] Derrick Nguyen and Bernard Widrow. The truck backer-upper: an example of self-learning in neural networks. *IJCNN*, 1989. 1, 3

[2] David Ha and Jürgen Schmidhuber. World models. *CoRR*, abs/1803.10122, 2018. URL http://arxiv.org/abs/1803.10122. 1, 3

[3] Yann LeCun. A path towards autonomous machine intelligence. OpenReview, 2022. Version 0.9.2, https://openreview.net/forum?id=BZ5a1r-kVsf. 1, 3

[4] Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie Shotton, and Gianluca Corrado. Gaia-1: A generative world model for autonomous driving. *ArXiv*, abs/2309.17080, 2023. URL https://api.semanticscholar.org/CorpusID:263310665. 1

[5] Danijar Hafner, Wilson Yan, and Timothy Lillicrap. Training agents inside of scalable world models. *arXiv preprint arXiv:2509.24527*, 2025. 1, 3

[6] Nicklas Hansen, Hao Su, and Xiaolong Wang. Learning massively multitask world models for continuous control. *ArXiv*, abs/2511.19584, 2025. URL https://api.semanticscholar.org/CorpusID:283251537. 1

[7] Lin Li, Qihang Zhang, Yiming Luo, Shuai Yang, Ruilin Wang, Fei Han, Mingrui Yu, Zelin Gao, Nan Xue, Xing Zhu, et al. Causal world modeling for robot control. *arXiv preprint arXiv:2601.21998*, 2026. 1

[8] Gaoyue Zhou, Hengkai Pan, Yann LeCun, and Lerrel Pinto. DINO-WM: world models on pre-trained visual features enable zero-shot planning. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu, editors, *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*, volume 267 of *Proceedings of Machine Learning Research*. PMLR / OpenReview.net, 2025. URL https://proceedings.mlr.press/v267/zhou25t.html. 1, 3, 4, 5, 7, 9, 10, 16, 18, 22

[9] Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Matthew Muckley, Ammar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zholus, et al. V-jepa 2: Self-supervised video models enable understanding, prediction and planning. *arXiv preprint arXiv:2506.09985*, 2025. 1, 3

[10] Amir Bar, Gaoyue Zhou, Danny Tran, Trevor Darrell, and Yann LeCun. Navigation world models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, pages 15791–15801. Computer Vision Foundation / IEEE, 2025. doi: 10.1109/CVPR52734.2025.01472. URL https://openaccess.thecvf.com/content/CVPR2025/html/Bar_Navigation_World_Models_CVPR_2025_paper.html. 1, 3, 21

[11] Raktim Gautam Goswami, Amir Bar, David Fan, Tsung-Yen Yang, Gaoyue Zhou, Prashanth Krishnamurthy, Michael Rabbat, Farshad Khorrami, and Yann LeCun. World models can leverage human videos for dexterous manipulation. *CoRR*, abs/2512.13644, 2025. doi: 10.48550/ARXIV.2512.13644. URL https://doi.org/10.48550/arXiv.2512.13644. 1, 3

[12] Delong Chen, Theo Moutakanni, Willy Chung, Yejin Bang, Ziwei Ji, Allen Bolourchi, and Pascale Fung. Planning with reasoning using vision language world model. *arXiv preprint arXiv:2509.02722*, 2025. 1

[13] Jürgen Schmidhuber. Learning to generate sub-goals for action sequences. In *Artificial neural networks*, pages 967–972, 1991. 1, 2

[14] Richard S. Sutton, Doina Precup, and Satinder Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artif. Intell.*, 112:181–211, 1999. URL https://api.semanticscholar.org/CorpusID:76564. 1, 2, 3

[15] Andrew G. Barto and Sridhar Mahadevan. Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems*, 13:41–77, 2003. URL https://api.semanticscholar.org/CorpusID:386824. 1, 2

[16] Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David Silver, and Koray Kavukcuoglu. Feudal networks for hierarchical reinforcement learning. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 3540–3549. PMLR, 2017. URL http://proceedings.mlr.press/v70/vezhnevets17a.html. 1, 3

[17] Ofir Nachum, Shixiang Shane Gu, Honglak Lee, and Sergey Levine. Data-efficient hierarchical reinforcement learning. In *Neural Information Processing Systems*, 2018. URL https://api.semanticscholar.org/CorpusID:43920938. 1, 3

[18] Seohong Park, Dibya Ghosh, Benjamin Eysenbach, and Sergey Levine. HIQL: offline goal-conditioned RL with latent states as actions. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/6d7c4a0727e089ed6cdd3151cbe8d8ba-Abstract-Conference.html. 1, 3, 18, 22

[19] Karl Pertsch, Youngwoon Lee, and Joseph J. Lim. Accelerating reinforcement learning with learned skill priors. In *Conference on Robot Learning*, 2020. URL https://api.semanticscholar.org/CorpusID:225039804. 1, 2, 3

[20] Abhishek Gupta, Vikash Kumar, Corey Lynch, Sergey Levine, and Karol Hausman. Relay policy learning: Solving long-horizon tasks via imitation and reinforcement learning. *ArXiv*, abs/1910.11956, 2019. URL https://api.semanticscholar.org/CorpusID:204955905. 1, 3

[21] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alexander Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario M Jauregui Ruano, Kyle Jeffrey, Sally Jesmonth, Nikhil Jayant Joshi, Ryan C. Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee, Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego M Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander Toshev, Vincent Vanhoucke, F. Xia, Ted Xiao, Peng Xu, Sichun Xu, and Mengyuan Yan. Do as i can, not as i say: Grounding language in robotic affordances. In *Conference on Robot Learning*, 2022. URL https://api.semanticscholar.org/CorpusID:247939706. 1

[22] Wenhao Li, Xiangfeng Wang, Bo Jin, and Hongyuan Zha. Hierarchical diffusion for offline decision making. In *International Conference on Machine Learning*, 2023. URL https://api.semanticscholar.org/CorpusID:260859191. 1, 3

[23] Chang Chen, Fei Deng, Kenji Kawaguchi, Caglar Gulcehre, and Sungjin Ahn. Simple hierarchical planning with diffusion. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL https://openreview.net/forum?id=kXHEBK9uAY. 1, 3, 18, 23

[24] Ce Hao, Anxing Xiao, Zhiwei Xue, and Harold Soh. Chd: Coupled hierarchical diffusion for long-horizon tasks. *ArXiv*, abs/2505.07261, 2025. URL https://api.semanticscholar.org/CorpusID:278502357. 1, 3

[25] Wancong Zhang, Basile Terver, Artem Zholus, Soham Chitnis, Harsh Sutaria, Mido Assran, Randall Balestriero, Amir Bar, Adrien Bardes, Yann LeCun, et al. Hierarchical planning with latent world models. *arXiv preprint arXiv:2604.03208*, 2026. 2, 3, 8, 10, 18

[26] Kevin Xie, Homanga Bharadhwaj, Danijar Hafner, Animesh Garg, and Florian Shkurti. Latent skill planning for exploration and transfer. In *International Conference on Learning Representations*, 2020. URL https://api.semanticscholar.org/CorpusID:233714270. 2, 3

[27] Christian Gumbsch, Noor Sajid, Georg Martius, and Martin Volker Butz. Learning hierarchical world models with adaptive temporal abstractions from discrete latent dynamics. In *International Conference on Learning Representations*, 2024. URL https://api.semanticscholar.org/CorpusID:268859646. 2, 3

[28] Soroush Nasiriany, Vitchyr H. Pong, Steven Lin, and Sergey Levine. Planning with goal-conditioned policies. In *Neural Information Processing Systems*, 2019. URL https://api.semanticscholar.org/CorpusID:202779996. 2, 3

[29] Kurtland Chua, Roberto Calandra, Rowan Thomas McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Neural Information Processing Systems*, 2018. URL https://api.semanticscholar.org/CorpusID:44177328. 2

[30] Michael Psenka, Michael Rabbat, Aditi Krishnapriyan, Yann LeCun, and Amir Bar. Parallel stochastic gradient-based planning for world models, 2026. URL https://arxiv.org/abs/2602.00475. 2, 3

[31] Junyoung Chung, Sungjin Ahn, and Yoshua Bengio. Hierarchical multiscale recurrent neural networks. *ArXiv*, abs/1609.01704, 2016. URL https://api.semanticscholar.org/CorpusID:1463401. 3

[32] Dan Haramati, Carl Qi, Tal Daniel, Amy Zhang, Aviv Tamar, and George D Konidaris. Hierarchical entity-centric reinforcement learning with factored subgoal diffusion. In *International Conference on Learning Representations*, volume 2026, pages 137484–137510, 2026. 3

[33] Suraj Nair and Chelsea Finn. Hierarchical foresight: Self-supervised learning of long-horizon tasks via visual subgoal generation. *ArXiv*, abs/1909.05829, 2019. URL https://api.semanticscholar.org/CorpusID:202565422. 3

[34] Lucy Xiaoyang Shi, Joseph J. Lim, and Youngwoon Lee. Skill-based model-based reinforcement learning. In Karen Liu, Dana Kulic, and Jeffrey Ichnowski, editors, *Conference on Robot Learning, CoRL 2022, 14-18 December 2022, Auckland, New Zealand*, volume 205 of *Proceedings of Machine Learning Research*, pages 2262–2272. PMLR, 2022. URL https://proceedings.mlr.press/v205/shi23a.html. 3

[35] Murtaza Dalal, Deepak Pathak, and Ruslan Salakhutdinov. Accelerating robotic reinforcement learning via parameterized action primitives. In *Neural Information Processing Systems*, 2021. URL https://api.semanticscholar.org/CorpusID:240070909. 3

[36] Jingwei Zhang, Jost Tobias Springenberg, Arunkumar Byravan, Leonard Hasenclever, Abbas Abdolmaleki, Dushyant Rao, Nicolas Manfred Otto Heess, and Martin A. Riedmiller. Leveraging jumpy models for planning and fast learning in robotic domains. *ArXiv*, abs/2302.12617, 2023. URL https://api.semanticscholar.org/CorpusID:257206001. 3

[37] Nicklas Hansen, Jyothir S. V, Vlad Sobal, Yann LeCun, Xiaolong Wang, and Hao Su. Hierarchical world models as visual whole-body humanoid controllers. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL https://openreview.net/forum?id=7wuJMvK639. 3

[38] Jinning Li, Chen Tang, Masayoshi Tomizuka, and Wei Zhan. Hierarchical planning through goal-conditioned offline reinforcement learning. *IEEE Robotics Autom. Lett.*, 7(4):10216–10223, 2022. doi: 10.1109/LRA.2022.3190100. URL https://doi.org/10.1109/LRA.2022.3190100. 3

[39] Jake Bruce, Michael D. Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, Yusuf Aytar, Sarah Bechtle, Feryal M. P. Behbahani, Stephanie C. Y. Chan, Nicolas Heess, Lucy Gonzalez, Simon Osindero, Sherjil Ozair, Scott E. Reed, Jingwei Zhang, Konrad Zolna, Jeff Clune, Nando de Freitas, Satinder Singh, and Tim Rocktäschel. Genie: Generative interactive environments. In *ICML*, 2024. 3

[40] Xiangyu Peng, Zangwei Zheng, Chenhui Shen, Tom Young, Xinying Guo, Binluo Wang, Hang Xu, Hongxin Liu, Mingyan Jiang, Wenjun Li, et al. Open-sora 2.0: Training a commercial-level video generation model in $200k. *arXiv preprint arXiv:2503.09642*, 2025. 3

[41] Thaddäus Wiedemer, Yuxuan Li, Paul Vicol, Shixiang Shane Gu, Nick Matarese, Kevin Swersky, Been Kim, Priyank Jaini, and Robert Geirhos. Video models are zero-shot learners and reasoners. *arXiv preprint arXiv:2509.20328*, 2025. 3

[42] Danijar Hafner, Timothy P. Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. In *ICLR*, 2020. 3

[43] Danijar Hafner, Timothy P. Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. In *ICLR*, 2021. 3

[44] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023. 3

[45] Nicklas Hansen, Xiaolong Wang, and Hao Su. Temporal difference learning for model predictive control, 2022. URL https://arxiv.org/abs/2203.04955. 3

[46] Nicklas Hansen, Hao Su, and Xiaolong Wang. TD-MPC2: scalable, robust world models for continuous control. In *ICLR*, 2024. 3

[47] Jan Robine, Marc Höftmann, Tobias Uelwer, and Stefan Harmeling. Transformer-based world models are happy with 100k interactions. In *ICLR*, 2023. 3

[48] Vincent Micheli, Eloi Alonso, and François Fleuret. Transformers are sample-efficient world models. In *ICLR*, 2023. 3

[49] Reuven Y Rubinstein. Optimization of computer simulation models with rare events. *European Journal of Operational Research*, 99(1):89–112, 1997. 3

[50] Uladzislau Sobal, Wancong Zhang, Kyunghyun Cho, Randall Balestriero, Tim GJ Rudner, and Yann LeCun. Learning from reward-free offline data: A case for planning with latent dynamics models. *Advances in Neural Information Processing Systems*, 38:43905–43941, 2026. 3

[51] Arjun Parthasarathy, Nimit Kalra, Rohun Agrawal, Yann LeCun, Oumayma Bounou, Pavel Izmailov, and Micah Goldblum. Closing the train-test gap in world models for gradient-based planning. *CoRR*, abs/2512.09929, 2025. doi: 10.48550/ARXIV.2512.09929. URL https://doi.org/10.48550/arXiv.2512.09929. 3

[52] Homanga Bharadhwaj, Kevin Xie, and Florian Shkurti. Model-predictive control via cross-entropy and gradient-based optimization. *ArXiv*, abs/2004.08763, 2020. URL https://api.semanticscholar.org/CorpusID:215827996. 3

[53] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis R. Bach and David M. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 2256–2265. JMLR.org, 2015. URL http://proceedings.mlr.press/v37/sohl-dickstein15.html. 3

[54] Jonathan Ho, Ajay Jain, and P. Abbeel. Denoising diffusion probabilistic models. *ArXiv*, abs/2006.11239, 2020. URL https://api.semanticscholar.org/CorpusID:219955663. 3, 4

[55] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL https://openreview.net/forum?id=St1giarCHLP. 3, 4

[56] James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, et al. Improving image generation with better captions. *Computer Science. https://cdn. openai. com/papers/dall-e-3. pdf*, 2(3):8, 2023. 3

[57] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, Seyedeh Sara Mahdavi, Raphael Gontijo Lopes, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. *ArXiv*, abs/2205.11487, 2022. URL https://api.semanticscholar.org/CorpusID:248986576. 3

[58] Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori Hashimoto. Diffusion-lm improves controllable text generation. *ArXiv*, abs/2205.14217, 2022. URL https://api.semanticscholar.org/CorpusID:249192356. 3

[59] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In Kostas E. Bekris, Kris Hauser, Sylvia L. Herbert, and Jingjin Yu, editors, *Robotics: Science and Systems XIX, Daegu, Republic of Korea, July 10-14, 2023*, 2023. doi: 10.15607/RSS.2023.XIX.026. URL https://doi.org/10.15607/RSS.2023.XIX.026. 3, 16, 18, 22

[60] Sudeep Dasari, Oier Mees, Sebastian Zhao, Mohan Kumar Srirama, and Sergey Levine. The ingredients for robotic diffusion transformers. In *IEEE International Conference on Robotics and Automation, ICRA 2025, Atlanta, GA, USA, May 19-23, 2025*, pages 15617–15625. IEEE, 2025. doi: 10.1109/ICRA55743.2025.11127493. URL https://doi.org/10.1109/ICRA55743.2025.11127493. 3

[61] Michael Janner, Yilun Du, Joshua B. Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 9902–9915. PMLR, 2022. URL https://proceedings.mlr.press/v162/janner22a.html. 3, 18, 23

[62] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua B. Tenenbaum, T. Jaakkola, and Pulkit Agrawal. Is conditional generative modeling all you need for decision-making? *ArXiv*, abs/2211.15657, 2022. URL https://api.semanticscholar.org/CorpusID:254044710. 3

[63] Haofei Lu, Dongqi Han, Yifei Shen, and Dongsheng Li. What makes a good diffusion planner for decision making? *ArXiv*, abs/2503.00535, 2025. URL https://api.semanticscholar.org/CorpusID:276741986. 3

[64] Po-Chen Ko, Jiayuan Mao, Yilun Du, Shao-Hua Sun, and Joshua B. Tenenbaum. Learning to act from actionless videos through dense correspondences. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL https://openreview.net/forum?id=Mhb5fpA1T0. 3, 10, 18

[65] Amelie Minji Kim, Anqi Wu, and Ye Zhao. Hierarchical diffusion motion planning with task-conditioned uncertainty-aware priors, 2025. URL https://arxiv.org/abs/2509.25685. 3

[66] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 10674–10685. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01042. URL https://doi.org/10.1109/CVPR52688.2022.01042. 10

[67] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4RL: datasets for deep data-driven reinforcement learning. *CoRR*, abs/2004.07219, 2020. URL https://arxiv.org/abs/2004.07219. 16

[68] Seohong Park, Kevin Frans, Benjamin Eysenbach, and Sergey Levine. Ogbench: Benchmarking offline goal-conditioned rl, 2025. URL https://arxiv.org/abs/2410.20092. 23

6 Appendix

6.1 Broader impact statement

This work studies hierarchical world models for goal-conditioned planning in navigation and manipulation. By decoupling state-space reachability from local control, our method enables long-horizon planning toward arbitrary goals without requiring task-specific reward functions, value functions, or pre-trained policies. In the long term, such task-agnostic planning frameworks could reduce the engineering effort and supervision required to deploy embodied agents in new environments, and could improve data efficiency by allowing a single learned world model to be reused across many downstream tasks. As with other advances in autonomous decision-making, deployment on physical systems should still involve careful testing, safety validation, and human oversight.

6.2 Reproducibility statement

To support reproducibility, we will release the datasets, training code, and planning code for **ID-WM** upon publication. We provide full architectural specifications, training hyperparameters, and planning hyperparameters for both the high-level and low-level world models in Appendix 6.14.1 (Tables 9, 10, 12, 11), as well as detailed environment descriptions and dataset generation procedures in Appendix 6.4. All baselines are trained on the same offline datasets and evaluated under identical task specifications, goal definitions, and environment interaction budgets as **ID-WM**, with baseline-specific hyperparameters reported in Appendix 6.7.

6.3 Large language models assistance disclosure

We used large language models to help refine wording and improve grammar in the manuscript. All technical claims, methodological choices, experimental designs, results, and conclusions were developed, executed, verified, and approved by the authors.

6.4 Environments and Dataset Generation

We evaluate **ID-WM** on three navigation and manipulation environment suites as shown in Figure 3. The environments are Wall, PushT, and Maze, with task horizon ranging from hundreds to thousands of environment steps.

a) **Wall:** Introduced in Zhou et al. [8], this is a simple 2D navigation environment that consists of two rooms separated by a wall with a door. The agent's goal is to navigate from a random starting location on one side of the wall to a random location on the other side, requiring the agent to pass through the door. We use the same dataset as in Zhou et al. [8], which consists of randomly generated 1920 trajectories of length 50.

b) **Maze:** This environment is introduced in Fu et al. [67], where a force-actuated 2-DoF ball navigates to an arbitrary location in a maze. In this work, we consider maze layout of four sizes, U (3×3), Medium (6×6), Large (7×10), and Giant (13×13), visualized in Figure 7 a). We generate 2000 random trajectories of length 100 for each maze layout.

c) **PushT:** Introduced by Chi et al. [59], this environment consists of a pusher agent and a T-shaped block. The objective is for both the agent and the block to reach specified target locations from random initial positions. We evaluate two variants of this task: in *Fix Goal*, the target state is a fixed configuration (the green T); in *Random Goal*, the target is randomly sampled from the environment's state space. Both settings are significantly more challenging than the formulation in Zhou et al. [8], where the goal is provided at a known, fixed horizon of 25 steps. In contrast, completing the task in our settings often requires hundreds of steps, rendering standard single-level planning infeasible. Following [8], we train on noisy, replayed trajectories. We further evaluate our method in a data-abundant regime by applying random translation and rotation augmentations to the existing dataset.

6.5 Task Visualizations

For the Wall, PushT, Giant Maze, Large Maze, Medium Maze, and U-Maze environments, we randomly sample 50 initial and goal states from each environment's state space (visualizations in

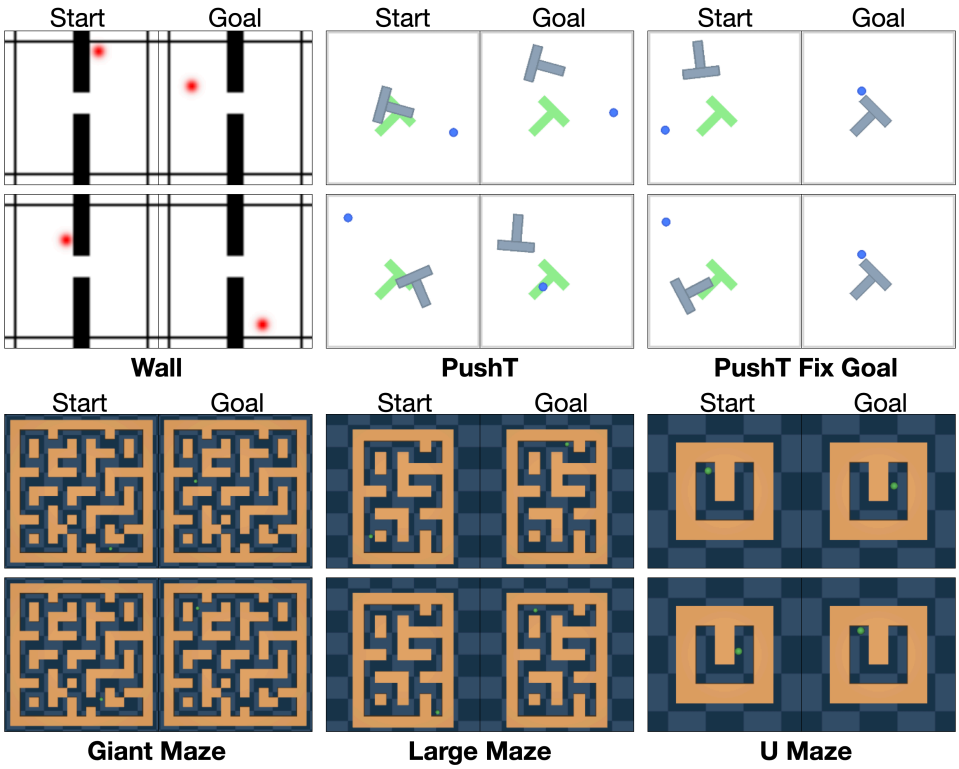


Figure 9: Examples of sampled initial and goal state pairs for the Wall, PushT, Giant Maze, Large Maze, and U-Maze environments

Figure 9). For PushT, the random-goal setting samples both the initial and goal states randomly from the environment, whereas the fixed-goal setting samples a random initial state while using the environment's fixed goal location. For the Wall environment, to ensure sufficient task difficulty, the initial and final object locations are constrained to lie on opposite sides of the wall. For the Giant Maze, Large Maze, Medium Maze, and U-Maze environments, we follow the random state sampling procedure used in OGBench: we sample arbitrary free cells for the ball and add continuous uniform noise around the corresponding grid-cell centers.

6.6 Data Augmentation

Data augmentation for PushT. For PushT we additionally apply geometric data augmentation to the offline dataset at each training step. The state vector contains end-effector position and velocity $(x_{\text{dot}}, y_{\text{dot}}, \dot{x}_{\text{dot}}, \dot{y}_{\text{dot}})$, T-block position (x_T, y_T), and T-block orientation $(\sin\theta_T, \cos\theta_T)$. One of three augmentation types is sampled per step:

- **Displacement** (40%): all positions shifted by a random integer offset $(\delta x, \delta y)$ chosen so that both the end-effector and the T-block remain within the screen boundaries $[50, 450]$ and $[100, 400]$ respectively.
- **Rotation** (30%): all positions and the T-block angle rotated by $\theta \sim \mathcal{U}(0, 2\pi)$ about the scene center; velocities and actions are rotated accordingly.
- **Displacement + Rotation** (30%): a half-range displacement combined with a small rotation $\theta \sim \mathcal{U}(0, 30^\circ)$.

6.7 Baselines

We compare **ID-WM** with five state-of-the-art offline learning methods including imitation learning (GCBC), offline RL (HIQL), diffusion planners (Diffuser, Hierarchical Diffuser), and planning with

world models (Single-level WM with MPC). For qualitative vision-based experiments, we compare **ID-WM** with AVDC [64].

a) **Single-level WM and Planning:** This baseline employs a flat world model based on Zhou et al. [8] to perform single-level planning via the Cross-Entropy Method (CEM) without hierarchical abstractions.

b) **Goal-conditioned Behavior Cloning (GCBC):** An imitation learning paradigm where a policy is trained to map the current state and a target goal observation directly to the expert action that leads toward that goal. For our implementation, we extend Diffusion Policy [59] by conditioning the policy on a future frame randomly sampled from the trajectory.

c) **HIQL [18]:** An offline reinforcement learning method that learns hierarchical representations by combining high-level subgoal-as-action policy with low-level subgoal-conditioned policy for long-horizon tasks.

d) **Diffuser [61]:** A trajectory-level diffusion model that performs planning by sampling entire state-action sequences. We replan after executing each action in the environment.

e) **Hierarchical Diffuser [23]:** An extension of Diffuser that decomposes the planning process into two temporal scales. A high-level module synthesizes intermediate subgoals, which are subsequently used by a low-level denoising network to generate a refined trajectory. We replan both the high level and the low level after executing each action in the environment.

f) **AVDC [64]:** This approach uses a video generation model to predict future states and then employs inverse kinematics (IK) to derive the necessary actions; we specifically use it to compare the quality of its generated video sequences against the high-level subgoals produced by our method.

g) **HWM [25]:** A concurrent hierarchical latent world model that increases the frameskip of the high-level model and plans over learned macro-actions. We compare with the DINO-WM variant of HWM. Since the official implementation of this variant is not open-sourced, we reimplement it following the reported settings and DINO-WM for unspecified details. We compare against HWM in Tables 2 and 3.

6.8 Hierarchical Planning with CEM

In this section, we provide the complete pseudo-code for our hierarchical planning procedure using CEM. Algorithm 2 expands upon the simplified overview presented in Algorithm 1, explicitly detailing the sampling and optimization loops for both high-level subgoal generation and low-level action execution.

6.9 Ablation: Effect of CEM Search Budget and Noise Regularization

We investigate the impact of the CEM search budget and the noise regularization strategy introduced in Section 3.2. By varying the number of CEM samples per level (maintaining a 10% elite fraction), we demonstrate that the performance of **ID-WM** scales effectively with test-time compute (Figure 10). Furthermore, planning with non-zero noise regularization ($\alpha > 0$) consistently yields higher success rates compared to the unregularized baseline. Unless otherwise noted, all reported results use the same planner configuration at both levels: 300 CEM samples, 30 elites, and a horizon of 5 (Table 7).

6.10 Ablation on Δ and number of subgoals

All results in the main paper use $\Delta = 5$ (25 environment steps with frameskip) and $H_{hi} = 5$ subgoals per high-level plan, without per-task tuning. State-based experiments use a 50-step noise schedule for both DDPM training and DDIM inference; vision-based experiments train with 250 steps and use 25 DDIM steps at inference.

To assess sensitivity to these choices, we sweep Δ and the number of subgoals on PushT-Aug with random goals (Table 4). Performance is robust for large Δ across a wide range of subgoal counts. For a fixed effective search horizon ($\Delta\times$ number of subgoals), splitting it evenly between the two levels is most effective: at an effective horizon of roughly 25, the success rate increases monotonically from 0.06 ($\Delta = 1$, 25 subgoals) to 0.96 ($\Delta = 5$, 5 subgoals). A very small Δ pushes the burden onto a long subgoal chain while contributing little per hop.

Algorithm 2 Hierarchical Planning with **ID-WM** (CEM)

Require: Initial state s_0, goal state s_g, high/low models f_{hi}, f_{lo}, costs C_{hi}, C_{lo}
Require: CEM parameters: N samples, M iterations, K elites, horizons H_{hi}, H_{lo}

$s_t \leftarrow s_0$
while goal s_g not reached **do**
  `# High-Level Plan (Subgoal Generation)`
  Initialize latent distribution parameters θ_{hi}
  for $m = 1$ to M **do**
    Sample N latent sequences $Z^{(n)} = (z_1^{(n)}, \ldots, z_{H_{hi}}^{(n)}) \sim \mathcal{D}(\theta_{hi})$
    for $n = 1$ to N **do**
      Rollout subgoals: $\hat{s}_t^{(n)} \sim f_{hi}(\hat{s}_{t-1}^{(n)}, z_t^{(n)})$ for $t \in [1, H_{hi}]$
      Compute trajectory cost $J_{hi}^{(n)} = C_{hi}(\hat{s}_{1:H_{hi}}^{(n)}, s_g)$
    end for
    Update θ_{hi} using top K elite sequences with the lowest J_{hi}
  end for
  Extract best subgoal sequence $(\hat{s}_1^*, \ldots, \hat{s}_{H_{hi}}^*)$
  `# Low-Level Plan (Action Generation)`
  for $i = 1$ to H_{hi} **do**
    Initialize action distribution parameters θ_{lo}
    for $m = 1$ to M **do**
      Sample N action sequences $A^{(n)} = (a_1^{(n)}, \ldots, a_{H_{lo}}^{(n)}) \sim \mathcal{D}(\theta_{lo})$
      for $n = 1$ to N **do**
        Rollout states: $\tilde{s}_t^{(n)} \sim f_{lo}(\tilde{s}_{t-1}^{(n)}, a_t^{(n)})$ for $t \in [1, H_{lo}]$
        Compute trajectory cost $J_{lo}^{(n)} = C_{lo}(\tilde{s}_{1:H_{lo}}^{(n)}, \hat{s}_i^*)$
      end for
      Update θ_{lo} using top K elite sequences with the lowest J_{lo}
    end for
    Extract best action sequence $(a_1^*, \ldots, a_{H_{lo}}^*)$
    `# Execute Planned Actions`
    Execute actions $(a_1^*, \ldots, a_{H_{lo}}^*)$ in the environment
    $s_t \leftarrow$ observe new state
  end for
end while

Table 4: Success rate on PushT-Aug (Random Goal) when sweeping the subgoal horizon Δ (rows) and the number of subgoals per high-level plan (columns). The default configuration ($\Delta = 5$, 5 subgoals) is underlined.

Δ \ #Subgoals	3	4	5	6	8	12	25
1	0.06	0.04	0.06	0.08	0.04	0.10	0.06
2	0.14	0.14	0.28	0.50	0.64	0.48	0.24
3	0.50	0.56	0.64	0.82	0.76	0.74	0.42
4	0.60	0.76	0.86	0.86	**1.00**	0.90	0.44
5	0.74	0.84	<u>0.96</u>	0.98	0.94	0.76	0.38

6.11 Steps-to-goal metric

Success rate measures whether a goal is reached but not how efficiently. To quantify trajectory efficiency, we report the average number of environment steps to success for all methods on PushT-Aug (Random Goal) and Medium Maze in Table 5, computed over successful trajectories only (N/A if no trajectory succeeded). **ID-WM** is the most efficient method on PushT-Aug and the most efficient among value-function-free methods on Medium Maze. HIQL's learned value function yields substantially shorter maze paths, reflecting a trade-off: the planning cost in Section 3.2 is the distance between the last imagined state and the goal, which optimizes for reaching the goal rather than for path efficiency. Within the same framework, a self-supervised value function could be substituted for this cost to improve path optimality.

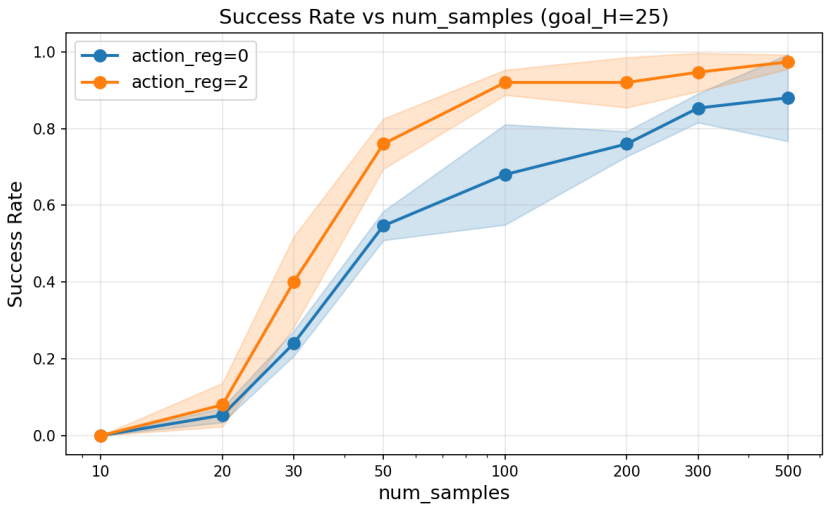


Figure 10: Success rate with respect to CEM samples and noise regularization.

Table 5: Average steps to success over successful trajectories (lower is better). N/A indicates that no trajectory succeeded.

Method	PushT-Aug (Rand.)	Medium Maze
HIQL	307.4	**34.1**
GCBC	1014.9	336.8
Diffuser	N/A	287.8
H-Diffuser	713.9	169.4
SL-WM	312.5	229.5
ID-WM	**234.7**	163.9

6.12 Test-time Compute

Table 6 reports the wall-clock inference time to reach one goal for each method in the state-based setting, measured on a single NVIDIA L40S GPU. GCBC and HIQL are reactive policies whose compute budget cannot be adapted at test time; GCBC is comparatively slow because of its Diffusion Policy backbone. Diffuser and H-Diffuser follow their original formulation of replanning a full trajectory after every action, which is the most expensive setting. Despite running CEM at both levels, **ID-WM** is roughly $27\times$ faster than Diffuser and $54\times$ faster than H-Diffuser, and faster than GCBC. The only faster methods are HIQL and SL-WM, both of which fail on long-horizon tasks.

The CEM search budget for **ID-WM** and SL-WM is identical in all reported results (Table 7): a horizon of 5 at each of our two levels gives an effective horizon of 25, matching SL-WM's flat horizon. **ID-WM** is slower only because of DDIM sampling in the high-level model.

Inference-time-matched SL-WM. Since SL-WM runs faster than **ID-WM** at the same CEM budget, we additionally match wall-clock time on PushT (Random Goal) by allowing SL-WM 45 replans ($3\times$). As shown in Table 8, tripling SL-WM's replanning budget yields only a marginal improvement (0.21 to 0.28) and remains far below **ID-WM** (0.77). This indicates that SL-WM's failure mode is compounding prediction error over long horizons rather than insufficient search, and that the gain comes from the hierarchy rather than from additional test-time compute.

6.13 Compute Resources

We train our state-based world models using 1 to 8 NVIDIA H200 Tensor Core GPUs, utilizing Distributed Data Parallel (DDP) to accelerate training when scaling to 8 GPUs. All baseline models are trained to convergence on a single NVIDIA H200 Tensor Core GPU. Planning evaluations are also conducted using a single NVIDIA H200 Tensor Core GPU. For the vision-based experiments,

Table 6: Inference time to reach one goal on a single NVIDIA L40S GPU.

Method	Inference time
HIQL	5.61 s
SL-WM	16.93 s
ID-WM	53.08 s
GCBC	73.3 s
Diffuser	24 min
H-Diffuser	48 min

Table 7: CEM planner configuration shared by **ID-WM** and SL-WM in all reported results.

Method	Samples	Elites	Replans	Horizon
SL-WM	300	30	15	25
ID-WM	300	30	15	5 (each level)

Table 8: Success rate on PushT (Random Goal) when matching either the CEM budget or the wall-clock inference time of SL-WM to **ID-WM**.

Method	Samples	Elites	Replans	Horizon	SR
SL-WM (CEM budget matched)	300	30	15	25	0.21
SL-WM (inference time matched)	300	30	45	25	0.28
ID-WM	300	30	15	5 (each level)	**0.77**

Table 9: Shared hyperparameters for training high-level dynamics model of **ID-WM** across all environments.

Name	Value
Emb Dim	56
Down Dims	[256]
Kernel	3
Groups	8
lr	1e−4
WDecay	0.1
Noise Timesteps	50
lr warmup steps	200

Table 10: Shared hyperparameters for planning with **ID-WM** across all environments.

Name	Value
Noise Reg. coefficient α	$\{0, 2\}$
τ	2
p	1
Max Replan	15
HL Horizon	5
LL Horizon	5
Top-k	30
Num Samples	300
Var Scale	1
Opt Steps	30

α is the coefficient on the noise regularization term $\mathcal{L}_{reg}$ in the high-level planning cost (Section 3.2); τ and p are its threshold and exponent.

we employ a CDiT [10] backbone and train the model using DDP across a cluster of 32 NVIDIA H200 Tensor Core GPUs.

6.14 Implementation Details

6.14.1 Implementations and Hyperparameters

Hyperparameters for **ID-WM** high-level dynamics model training and planning are listed in Table 9 and Table 10, respectively. Hyperparameters for low-level WM training are listed in Table 11 and Table 12, while those for low-level WM planning are also included in Table 10. The high-level dynamics model and low-level WM share the same CEM planner configurations, except that action regularization is applied only in high-level planning. We use consistent hyperparameters across tasks, except for the observation history length in the Wall environment, which follows DINO-WM.

For success metrics, we continue using all state-based evaluation metrics from DINO-WM. Success is evaluated, and the planner may stop, after each plan is fully executed.

6.14.2 Baseline Implementations and Hyperparameters

Unless otherwise stated, all baselines are trained on the same offline datasets and evaluated using the same task specifications, goal definitions, and environment interaction budget as **ID-WM**.

Table 11: Environment-dependent hyperparameters for WM training. T_o denotes the observation history length.

	T_o	Frameskip	Dataset Size	Traj. Len.
PushT	3	5	18500	100–300
Wall	1	5	1920	50
U-Maze	3	5	2000	100
Medium Maze	3	5	2000	100
Large Maze	3	5	2000	100
Giant Maze	3	5	2000	100

Table 12: Shared hyperparameters for WM training.

Name	Value
Predictor lr	5e−4
Action encoder lr	5e−4
Depth	6
Heads	8
MLP dim	64
Dropout	0
Emb. dropout	0
Pool	Mean

Table 13: Environment-dependent hyperparameters for GCBC training. T_o: observation history length; Act. Window: minimum timestep gap between current state and goal state; Pred. Horizon: action length generated per diffusion forward pass; Act. Steps: actions executed before re-planning.

	T_o	Act. Window	Pred. Horizon	Act. Steps
PushT	2	9	10	8
Wall	2	19	20	20
U-Maze	2	19	20	20
Medium Maze	2	19	20	20
Large Maze	2	19	20	20
Giant Maze	2	19	20	20

Table 14: Shared hyperparameters for GCBC training.

Name	Value
lr	1e−4
Weight decay	0
Diffusion steps	100
Layers	8
Heads	4
Embedding dim	256
Attn. dropout	0.1
Emb. dropout	0

6.14.3 Single-level WM and Planning

We implement a single-level WM with transformer-based predictor following DINO-WM [8]. The training and inference details are identical to those of the low level WM in Section 3.1.2 with implementation details in Section 6.14.1. During inference, the planner optimizes planning cost $\|\hat{s}_{H_{lo}} - s^*\|_2$. For state-based WM, we adapt DINO-WM and remove the frozen pre-trained DINOv2 encoder. This single-level WM baseline isolates the effect of hierarchy and tests whether subgoal proposal mitigates cumulative prediction error.

6.14.4 Goal-conditioned Behavior Cloning (GCBC)

GCBC trains a policy to imitate dataset actions conditioned on both the current state and a desired goal state. For training, given an offline trajectory $\{(s_t, a_t)\}_{t=0}^{T}$, we randomly sample a current state s_t and a future state s_{t+k} where $k > 0$. The policy is trained to predict dataset action a_t with the training objective:

$$\mathcal{L}_{\text{GCBC}} = \mathbb{E}_{(s_t, a_t, g)\sim\mathcal{D}} \left[\|\pi_\theta(s_t, g) - a_t\|_2^2 \right], \tag{7}$$

Specifically, we implement GCBC as a Diffusion Policy [59] conditioned on the concatenation of the current observation and goal state to predict an action chunk.

We include GCBC as a representative of direct goal-conditioned imitation learning and as a baseline for what behavioral cloning alone can achieve. Hyperparameters are listed in Table 13 and Table 14.

6.14.5 HIQL [18]

HIQL decomposes offline goal-conditioned RL into two levels of hierarchy. A high-level policy proposes a subgoal representation $\phi([s; w])$, where w is an intermediate waypoint state sampled from the same trajectory, and a low-level policy executes primitive actions to reach that subgoal. Both policies are trained offline via Advantage-Weighted Regression (AWR) using a shared implicit value function $V_\theta(s, g)$.

The value function is trained with an IQL-style expectile regression objective:

$$\mathcal{L}_V(\theta_V) = \mathbb{E}_{(s,s')\sim\mathcal{D},\, g\sim p(g|\tau)} \left[L_2^\tau \left(r(s, g) + \gamma V_{\bar{\theta}_V}(s', g) - V_{\theta_V}(s, g) \right) \right].$$

where L_2^τ denotes the asymmetric squared expectile loss with parameter τ. Given the shared value function, both policies are updated via AWR. For the low-level policy π^l, the advantage measures

Table 15: Shared hyperparameters for HIQL across all evaluation environments.

Name	Value
lr	3e−4
Discount factor γ	0.99
Expectile τ	0.7
High-level α	3.0
Low-level α	3.0
Subgoal steps k_{sub}	25
Actor/value hidden dims	$(512, 512, 512)$
Train steps	10^7

Table 16: Environment-dependent hyperparameters for Diffuser and HDiffuser. H: planning horizon; J: subsampling jump (HDiffuser high-level planner only).

	Diffuser	HDiffuser	
	H	H	J
PushT	128	125	25
Wall	48	45	15
U-Maze	48	75	25
Medium Maze	48	75	25
Large Maze	48	75	25
Giant Maze	96	75	25

Table 17: Shared hyperparameters for Diffuser and HDiffuser training.

Name	Value
lr	2e−4
Diffusion steps T	256
EMA decay	0.995
U-Net dim. mults	$(1, 4, 8)$

progress toward subgoal w. For the high-level policy π^h, the advantage measures progress toward the final goal g upon reaching subgoal w.

We use the HIQL implementation from the OGBench repository [68]. In all state-based experiments, no visual encoder is used, and ϕ processes raw state vectors directly. As HIQL is stochastic, all results are averaged over 3 random seeds. Hyperparameters are listed in Table 15.

6.14.6 Diffuser [61]

Diffuser frames offline goal-conditioned planning as a trajectory denoising problem. A temporal U-Net processes a full trajectory $\tau = (s_0, a_0, \ldots, s_{H-1}, a_{H-1})$ of planning horizon H, and iterative reverse diffusion over T steps is guided by goal conditioning to produce trajectories that reach a specified goal state. At inference, the current state is inpainted at the start of the trajectory and the goal state at the end; the remaining timesteps are filled by denoising. All environments share the same base configuration. Only the planning horizon H is adjusted (in Table 16).

Detailed implementation choices for inference are expanded in Section 6.14.8.

6.14.7 Hierarchical Diffuser [23]

H-Diffuser [23] extends Diffuser with a two-level hierarchy. A high-level diffusion model plans jumpy subgoals spaced J timesteps apart, producing a compressed sequence of length $\lceil H_{\text{hl}}/J \rceil$ that covers a larger temporal horizon at lower computational cost. A *low-level* diffusion model then plans a dense trajectory of length H_{ll} between consecutive subgoals, conditioned on the current state and the subgoal as the terminal constraint.

Hyperparameters shared with Diffuser are listed in Table 17. Environment-specific hyperparameters are listed in Table 16.

Detailed implementation choices for inference are expanded in Section 6.14.8.

6.14.8 Baseline Tuning for Diffuser & H-Diffuser

The original Diffuser paper uses two execution strategies depending on the environment: for D4RL locomotion tasks, only the first action of each plan is executed and the model replans at every step

(**execute one action + replan**); for maze2d and Kuka block-stacking, the full planned trajectory is executed open-loop without replanning (**execute plan**). H-Diffuser adopts the same two strategies at the subgoal level. To enable a fair comparison with our method, we additionally evaluate both models under a third strategy in which the full plan is executed and the model then replans (**execute plan + replan**). Results under all three strategies are reported in Tables 18 and 19.

Table 18: Diffuser results across execution settings and PushT variants.

Model	Execution	PushT				PushT-Aug			
		Fix Goal		Random Goal		Fix Goal		Random Goal	
		SR	Cov	SR	Cov	SR	Cov	SR	Cov
Diffuser	execute plan + replan	0.00	0.03	0.00	0.03	0.28	0.37	0.06	0.19
Diffuser	execute one action + replan	0.00	0.11	0.00	0.06	0.04	0.27	0.02	0.15
Diffuser	execute plan	0.00	0.05	0.00	0.03	0.00	0.09	0.00	0.13

Table 19: H-Diffuser results across execution settings and PushT variants.

Model	Execution	PushT				PushT-Aug			
		Fix Goal		Random Goal		Fix Goal		Random Goal	
		SR	Cov	SR	Cov	SR	Cov	SR	Cov
H-Diffuser	execute plan + replan	0.38	0.28	0.00	0.08	0.56	0.51	0.46	0.36
H-Diffuser	execute one action + replan	0.00	0.25	0.02	0.11	0.56	0.56	0.26	0.50
H-Diffuser	execute plan	0.04	0.09	0.00	0.06	0.02	0.17	0.00	0.10

6.15 High-Level World Model Training Convergence

We further analyze the training-time convergence behavior of the high-level world model, extending the qualitative study in Figure 4. While Figure 4 visualizes the final learned neighboring-state distribution, Figure 11 shows how this distribution emerges over the course of training.

Specifically, we fix validation observations s_0 and periodically sample predicted future states from intermediate checkpoints of the high-level WM. The first row shows the initial observation s_0 sampled from the validation set. The lower rows visualize generated future states $\hat{s}_\Delta = f_{hi}(s_0, z, \Delta)$ for horizons $\Delta \in \{1, 3, 5\}$, where each prediction is obtained by sampling latent noise vectors $z \sim \mathcal{N}(0, I)$. Each column corresponds to a different number of training steps on a logarithmic scale: 20, 60, 200, 600, 2000, 6000, and 20000.

As training begins, the predicted distributions are highly dispersed and often fail to preserve the structure of the PushT scene. In early checkpoints, the sampled futures do not consistently respect the local geometry of the object or the physically plausible displacement induced by the agent-object interaction. As training progresses, the generated samples become increasingly concentrated around coherent future configurations. By approximately 200 to 600 training steps, the model begins to capture the coarse reachable region for each temporal horizon, and later checkpoints further sharpen the predicted distribution while preserving meaningful sample diversity.

6.16 Comparison of ID-WM and naive WM with equivalent frameskip

We visualize the difference between our proposed **ID-WM** high-level world model and a naive WM without implicit dynamics. The left column shows the sampled noise/action distribution. In Figure 12, the middle column shows the neighboring states predicted by **ID-WM**, and the right column shows predictions from a naive WM with frameskip matched to the **ID-WM** subgoal horizon. **ID-WM** yields a smooth and locally consistent reachable set. The naive WM produces irregular mappings: nearby samples can jump to distant states, leave the local reachable neighborhood, or form discontinuous regions. This indicates that long-skip prediction with a naive world model is prone to out-of-distribution subgoals and discontinuous action-to-state geometry, making it less suitable for planning.

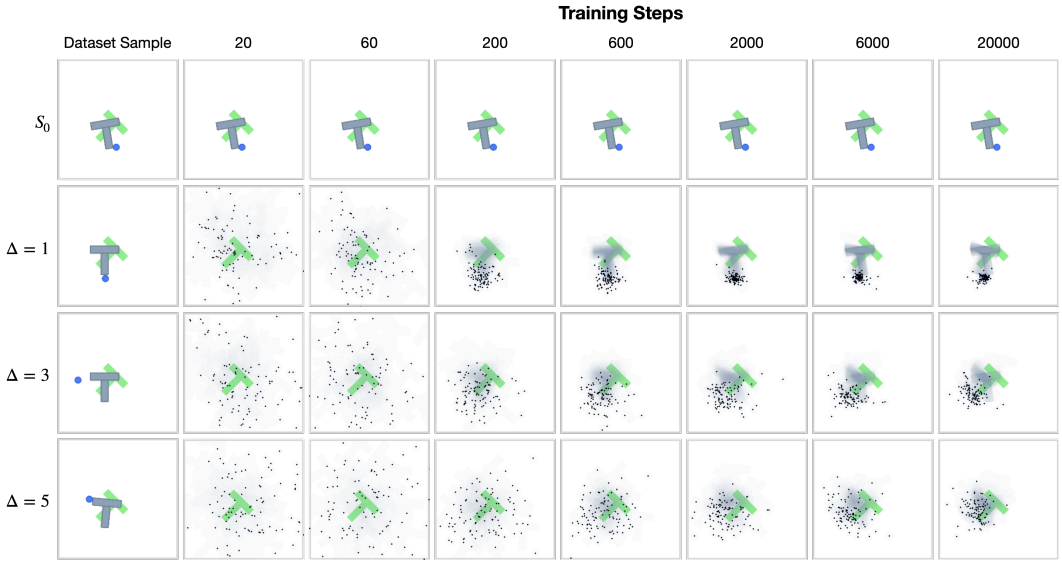


Figure 11: Training-time convergence of predicted future-state distributions for the PushT environment

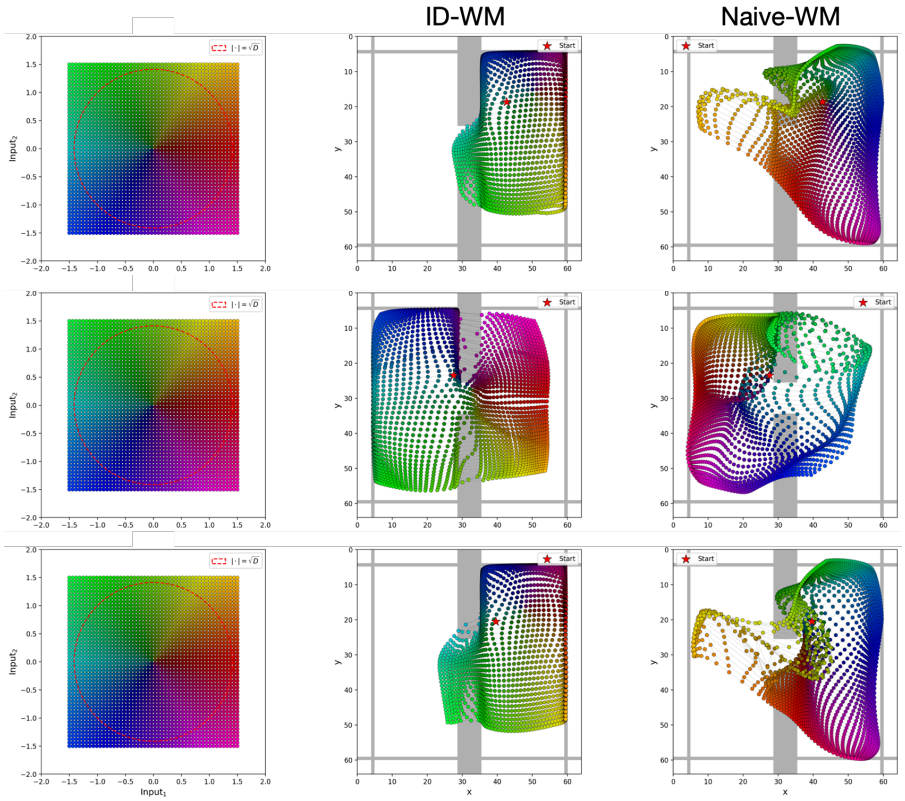


Figure 12: Visualization of the induced state mappings by **ID-WM** and naive WM.

NeurIPS Paper Checklist

1. **Claims**

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

   Answer: [Yes]

   Justification: We show the effectiveness of **ID-WM** both through qualitative subgoal and plan visualizations as well as quantitative task-planning performance for unseen initial position and goal pairs.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

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

   Answer: [Yes]

   Justification: We include a limitation section in the Conclusion section.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

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

Answer: [N/A]

Justification: No theory results are included in this work.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

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

   Answer: [Yes]

   Justification: All code and data will be publicly available at `https://id-wm.github.io` to ensure reproducibility.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

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

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

Answer: [Yes]

Justification: All code and data will be publicly available at https://id-wm.github.io to ensure reproducibility.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

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

Answer: [Yes]

Justification: Experimental setting for all experiments was included in the main text and in the appendix, including hyperparameters for ours and baselines.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

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

Answer: [Yes]

Justification: Experiments are conducted on three random seeds, and standard deviation bars have been provided.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

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

   Answer: [Yes]

   Justification: Compute resources has been provided in Section 6.13.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

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

   Answer: [Yes]

   Justification: This work focuses on hierarchical world models and planning, conforms with the NeurIPS Code of Ethics

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

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

    Answer: [Yes]

    Justification: This is provided in Section 6.1.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

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

    Answer: [N/A]

    Justification: This work doesn't include any sensitive data.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

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

    Answer: [Yes]

    Justification: All datasets and environment used in the paper has been cited.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

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

    Answer: [Yes]

    Justification: Full details of our data generation has been included in the appendix.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

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

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

    Answer: [N/A]

    Justification: This work doesn't contain any human subject.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

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

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

    Answer: [N/A]

    Justification: No study participants are involved in this work.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

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

Answer: [Yes]

Justification: This is included in Section 6.3

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.