

OGBENCH: BENCHMARKING OFFLINE GOAL-CONDITIONED RL

Anonymous authors

Paper under double-blind review

ABSTRACT

Offline goal-conditioned reinforcement learning (GCRL) is a major problem in reinforcement learning (RL) because it provides a simple, unsupervised, and domain-agnostic way to acquire diverse behaviors and representations from unlabeled data without rewards. Despite the importance of this setting, we lack a standard benchmark that can systematically evaluate the capabilities of offline GCRL algorithms. In this work, we propose OGBench, a new, high-quality benchmark for algorithms research in offline goal-conditioned RL. OGBench consists of 7 types of environments, 59 datasets, and reference implementations of 6 representative offline GCRL algorithms. We have designed these challenging and realistic environments and datasets to directly probe different capabilities of algorithms, such as stitching, long-horizon reasoning, and the ability to handle high-dimensional inputs and stochasticity. While representative algorithms may rank similarly on prior benchmarks, our experiments reveal stark strengths and weaknesses in these different capabilities, providing a strong foundation for building new algorithms. Videos: <https://ogbenchauthors.github.io/ogbench-anon/>

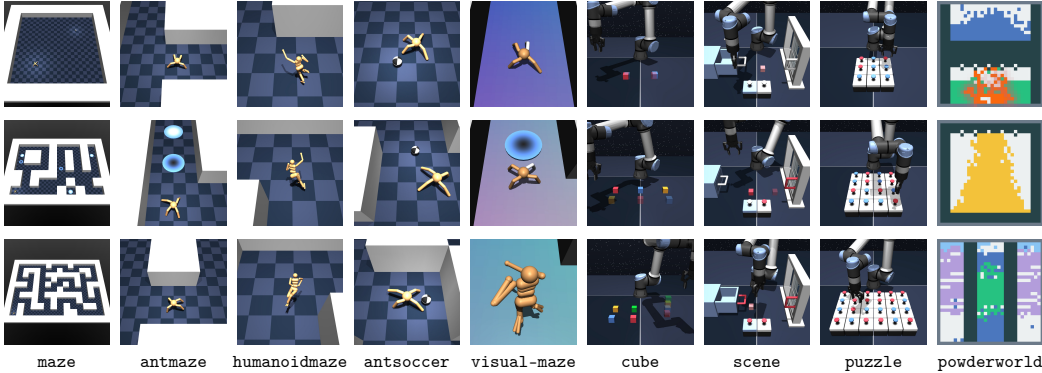


Figure 1: **OGBench Overview.** OGBench provides a variety of state- and pixel-based locomotion, manipulation, and drawing tasks that are designed to exercise diverse challenges in offline goal-conditioned RL, such as stitching, long-horizon reasoning, and stochastic control.

1 MOTIVATION

Why offline goal-conditioned reinforcement learning (RL)? The enduring trend in modern machine learning is to simplify domain-specific assumptions and scale up the data. In computer vision and natural language processing, the strongest general-purpose models are trained via simple *unsupervised* objectives on raw, unlabeled data, such as next-token prediction, contrastive learning, and masked auto-encoding. What analogous paradigm could enable data-driven unsupervised learning for reinforcement learning? Ideally, such a framework should be able to produce from data a generalist policy that can be directly queried or adapted to solve a variety of downstream tasks, much like how generative language models trained via next-token prediction can be easily adapted to everyday tasks.

We posit that a natural analogy to data-driven unsupervised learning in RL is **offline goal-conditioned RL (GCRL)**. The objective of offline goal-conditioned RL is fully unsupervised, remarkably simple, and requires no domain knowledge: it merely aims to learn to reach any state from any other state in the dataset in the fewest number of steps. However, mastering this simple objective is exceptionally difficult: the agent needs to not only acquire diverse skills to efficiently navigate the state space, but also have a deep, complete understanding of the underlying world and dataset. As a result, offline

goal-conditioned RL yields a highly capable general-purpose multi-task policy as well as rich, useful representations that can be adapted to solve a variety of downstream tasks (Ghosh et al., 2023; Kim et al., 2024). Indeed, for such simplicity and generality, interest in (offline) goal-conditioned RL has recently surged to the extent that even a standalone workshop on goal-conditioned RL was held at a machine learning conference.¹

Why a new benchmark? Despite the importance of and increasing interest in offline goal-conditioned RL, we currently lack a standard benchmark that can systematically assess the capabilities of offline GCRL algorithms, such as the ability to stitch, perform long-horizon reasoning, and handle stochasticity. Prior works in offline goal-conditioned RL (Eysenbach et al., 2022; Ma et al., 2022; Park et al., 2023; Myers et al., 2024) have mainly used either existing datasets for standard offline RL tasks without modification (e.g., D4RL (Fu et al., 2020)), relatively simple goal-conditioned tasks (e.g., Fetch (Plappert et al., 2018)), or tasks tailored to demonstrate the individual abilities of the proposed methods. This often results in limited evaluation. Prior works often evaluate their multi-task policies only on a single task (when using datasets not originally designed for offline GCRL), or learn relatively simple behaviors. While there exist some prior tasks tailored to evaluate *individual* properties of offline GCRL such as stitching or generalization (Yang et al., 2023; Ghugare et al., 2024), we lack a comprehensive, standardized benchmark that exhaustively assesses various properties of offline GCRL algorithms with diverse, challenging tasks.

Therefore, we introduce a benchmark named **Offline Goal-Conditioned RL Benchmark (OGBench)** in this work (Figure 1). The primary goals of this benchmark are to facilitate *algorithms research* in offline goal-conditioned RL and to provide a set of complex tasks that can unlock the potential of offline GCRL. Our benchmark introduces 7 types of environments and 59 datasets across robotic locomotion, robotic manipulation, and drawing, and provides well-tuned reference implementations of 6 representative offline GCRL methods. These datasets, tasks, and implementations are carefully designed. Tasks are designed in a way that complex behaviors can naturally emerge when successfully solved, and that they pose diverse algorithmic challenges in offline GCRL, such as goal stitching, stochastic control, long-horizon reasoning, and more. Dataset and task difficulties are carefully adjusted to highlight stark contrasts between algorithms across multiple criteria. The entire benchmark is designed to minimize unnecessary computational overhead and maximize usability such that any researcher can easily iterate and evaluate new ideas. We believe OGBench serves as a solid foundation for developing ideal algorithms for goal-conditioned and unsupervised RL from data.

2 PROBLEM SETTING

The offline goal-conditioned RL problem is defined by a controlled Markov process $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mu, p)$ (i.e., a Markov decision process (MDP) without rewards) and an unlabeled dataset \mathcal{D} , where \mathcal{S} denotes the state space, \mathcal{A} denotes the action space, $\mu(s) \in \Delta(\mathcal{S})^2$ denotes the initial state distribution, and $p(s' | s, a) : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ denotes the transition dynamics function. $\Delta(\mathcal{X})$ denotes the set of probability distributions defined on a set \mathcal{X} . The dataset $\mathcal{D} = \{\tau^{(n)}\}_{n \in \{1, 2, \dots, N\}}$ consists of unlabeled trajectories $\tau^{(n)} = (s_0^{(n)}, a_0^{(n)}, s_1^{(n)}, \dots, s_{T_n}^{(n)})$.

The objective of offline goal-conditioned RL is to learn to reach any state from any other state in the minimum number of time steps. Formally, offline GCRL aims to learn a goal-conditioned policy $\pi(a | s, g) : \mathcal{S} \times \mathcal{S} \rightarrow \Delta(\mathcal{A})$ that maximizes the objective $\mathbb{E}_{\tau \sim p(\tau | g)}[\sum_{t=0}^T \gamma^t \delta_g(s_t)]$ for all $g \in \mathcal{S}$, where $T \in \mathbb{N}$ denotes the episode horizon, $\gamma \in (0, 1)$ denotes the discount factor, $p(\tau | g)$ denotes the distribution given by $p(\tau | g) = \mu(s_0)\pi(a_0 | s_0, g)p(s_1 | s_0, a_0) \cdots p(s_T | s_{T-1}, a_{T-1})$, and $\delta_g(\cdot)$ denotes the Dirac delta “function”³ at g . Note that we use the entire state space as the goal space (i.e., a goal is simply a full state, not part of the state like only the x - y position of the agent). This choice makes the objective fully unsupervised, making it suitable for domain-agnostic training from unlabeled data. Our goal in this paper is to propose a new benchmark in offline GCRL. That is, formally speaking, to specify the dynamics and dataset for each task we introduce.

¹<https://goal-conditioned-rl.github.io/2023/>

²We denote placeholder variables in gray throughout the paper.

³In a discrete MDP, $\delta_g(s)$ is equal to an indicator function $\mathbf{1}_{\{g\}}(s)$. In a continuous MDP, it is technically not well-defined in its current form. It can be made precise using measure-theoretic notation or the distribution theory, but we choose to avoid them for simplicity.

Table 1: **Properties of benchmark tasks.** We summarize the properties of benchmark tasks commonly used in prior works (above the line) and our OGBench tasks (below the line).

Benchmark Task	Type ¹	Longest Task ²	# Subtasks ³	Test Stitching? ⁴	Have Stoch. Tasks? ⁵	Support Pixels? ⁶	Multi-Goal? ⁷	Dependency ⁸
D4RL AntMaze	Loco.	≈ 400	-	×	×	×	×	MuJoCo
D4RL Kitchen	Manip.	≈ 250	4	○	×	×	×	MuJoCo
Roboverse	Manip.	≈ 100	1-2	○	×	○	○	PyBullet
Fetch	Manip.	≈ 50	1	×	×	×	○	MuJoCo
AntMaze (ours)	Loco.	≈ 1000	-	○	○	○	○	MuJoCo
HumanoidMaze (ours)	Loco.	≈ 3000	-	○	×	○	○	MuJoCo
AntSoccer (ours)	Loco.	≈ 1000	-	○	×	×	○	MuJoCo
Cube (ours)	Manip.	≈ 400	1-4	○	×	○	○	MuJoCo
Scene (ours)	Manip.	≈ 250	2-7	○	×	○	○	MuJoCo
Puzzle (ours)	Manip.	≈ 600	2-24	○	×	○	○	MuJoCo
Powderworld (ours)	Draw.	≈ 100	-	○	○	○	○	PyTorch

¹ Environment type (locomotion, manipulation, or drawing). ² The (approximate) minimum number of environment steps to solve the longest task.

³ The number of atomic behaviors (e.g., pick-and-place) in each manipulation task. ⁴ Does it contain tasks that require goal stitching?

⁵ Does it contain tasks with stochastic dynamics? ⁶ Does it support pixel-based observations? ⁷ Does it use multiple goals for evaluation?

⁸ The main dependency of the benchmark.

3 HOW HAVE PRIOR WORKS BENCHMARKED OFFLINE GCRL?

While many excellent offline GCRL algorithms have been proposed so far, the community currently lacks a standardized way to evaluate their performance, unlike other fields in RL (Brockman et al., 2016; Tassa et al., 2018; Fu et al., 2020; Terry et al., 2021). Tasks used by prior works in offline GCRL are often limited for providing a proper evaluation for various reasons. For example, many works directly use tasks in existing offline RL benchmarks (not necessarily designed for offline goal-conditioned RL), such as D4RL AntMaze and Kitchen (Fu et al., 2020). However, since the tasks are designed for single-task offline RL, they often evaluate their multi-task policies only on the single, original goal when using these tasks (Eysenbach et al., 2022; Park et al., 2023; Zheng et al., 2024b; Myers et al., 2024), which results in limited evaluation. Some employ online GCRL tasks provided by Plappert et al. (2018) (e.g., Fetch tasks) with policy-collected datasets (Ma et al., 2022; Yang et al., 2023), but these tasks are mostly atomic (e.g., single pick-and-place) and do not sufficiently address various challenges in offline GCRL, such as long-horizon reasoning and goal stitching. While Roboverse (Fang et al., 2022; Zheng et al., 2024a) provides pixel-based manipulation tasks, the tasks are still relatively atomic and it requires installing multiple, fragmented dependencies, making it less approachable for researchers. Some prior works construct bespoke tasks and datasets to individually study the important and specific features of their algorithms (e.g., stitching (Ghugare et al., 2024) and stochasticity (Ma et al., 2022; Myers et al., 2024; Wang et al., 2024)), but it is often not entirely clear how algorithms compare to one another or if the new capabilities of new methods (e.g., stitching) come at the loss of other capabilities. These limitations of previous evaluation tasks have motivated us to create a new benchmark. In this work, we introduce a set of diverse tasks that cover various challenges in offline GCRL, enabling a much more thorough and multi-faceted evaluation than previous tasks. We summarize the properties of the previous tasks and our new tasks in Table 1 and refer to Appendix C for further discussion on related work.

4 OVERVIEW OF OGBENCH

We now introduce our benchmark, **Offline Goal-Conditioned RL Benchmark (OGBench)**. The primary goal of this benchmark is to provide tasks and datasets to unlock the full potential of offline goal-conditioned RL. To this end, we pose diverse challenges in offline GCRL throughout the benchmark in such a way that researchers can easily test and iterate on algorithmic ideas, and that complex, intriguing behaviors can naturally emerge when successful. OGBench consists of 7 types of environments, 59 datasets, and reference implementations of 6 representative offline GCRL algorithms. In the following sections, we first describe the challenges in offline GCRL (Section 5) and then outline our core design philosophies (Section 6). We next introduce the tasks and datasets (Section 7) and present the benchmarking results of the current algorithms (Section 8).

5 CHALLENGES IN OFFLINE GOAL-CONDITIONED RL

Offline goal-conditioned RL, despite its simplicity, is a challenging problem. Here, we discuss the major challenges in offline GCRL, which will motivate the design choices in our benchmark tasks.

(1) Learning from suboptimal, unstructured data: An ideal offline GCRL algorithm should be able to learn an effective multi-task policy from diverse and suboptimal data. This is especially important, considering that suboptimal (yet diverse) data is much cheaper to collect than curated, expert datasets (Lynch et al., 2019), and that the very use of large, diverse, unstructured data is one of

the foundations for the success of modern machine learning. Reflecting this challenge, we provide datasets with high diversity and varying suboptimality in this benchmark to challenge the capabilities of offline GCRL algorithms.

(2) Goal stitching: Another important challenge is to stitch the initial and final states of different trajectories to learn more diverse behaviors. We call this “goal stitching.” Goal stitching is different from “regular” stitching in offline RL, which applies only when the dataset is suboptimal. Unlike regular stitching, goal stitching applies *even when the dataset only consists of optimal, expert trajectories*, because we can often acquire more diverse goal-reaching behaviors by stitching multiple trajectories together, regardless of their optimality. For instance, an agent can stitch two atomic pick-and-place behaviors to sequentially move two objects in a single episode, even when the dataset does not contain any double pick-and-place behaviors. Goal stitching is crucial to learning diverse behaviors in many real-world applications with high behavioral diversity and large state spaces. In our benchmark, we introduce many tasks with large state spaces to assess the ability to stitch goals.

(3) Long-horizon reasoning: Long-horizon reasoning refers to the capability of navigating from a starting state to a goal state that is many steps apart. This challenge is important in many real-world tasks like autonomous driving or assembly, which may require several hours of continuous control or achieving dozens of subtasks. To substantially challenge the long-horizon reasoning ability of offline GCRL methods, we introduce tasks that are more than 5 times longer than previously used ones in terms of both the episode length and the number of subtasks (Table 1).

(4) Handling stochasticity: Another prominent challenge in offline GCRL is the ability to deal with stochastic environments. Correctly handling stochasticity is very important in practice, because virtually any real-world environment is stochastic due to partial observability. Yet, many works in offline GCRL assume *deterministic* dynamics to exploit the metric structure of temporal distances (Wang et al., 2023) or to enable hierarchical control (Park et al., 2023), at the cost of being optimistically biased in stochastic environments (Wang et al., 2023; Park et al., 2023). Correctly handling environment stochasticity while fully exploiting the recursive subgoal structure of GCRL remains an open problem. Since most previous tasks used to evaluate offline GCRL methods have deterministic dynamics (Table 1), we introduce several challenging tasks with stochastic dynamics in this benchmark.

6 DESIGN PRINCIPLES

Next, we discuss the design principles underlying our benchmark tasks. The tasks are intended to exercise the major challenges in offline GCRL in the previous section, while providing a set of high-quality tasks that provide not only a toolkit for algorithms research and evaluation, but also a platform for vividly illustrating the capabilities of offline GCRL with compelling and complex domains.

(1) Realistic and exciting tasks: The tasks should be realistic yet exciting enough while posing diverse challenges in offline goal-conditioned RL. Imagine a robot arm watching random movements of a puzzle and then solving it zero-shot at test time, a humanoid robot navigating through a labyrinth, or an agent painting cool pictures using different types of brushes. In this benchmark, we design new tasks such that these kinds of exciting behaviors can naturally emerge when an RL agent properly stitches different trajectory segments together (up to 24; see Table 1) or successfully generalizes. At the same time, we make sure our tasks exhaustively cover major challenges in offline goal-conditioned RL, such as long-horizon reasoning, stochastic control, and goal stitching (Section 5).

(2) Appropriate difficulty: The tasks and datasets should have appropriate levels of difficulty to properly evaluate different algorithms. In other words, they should be of high quality *for benchmarking*. No matter how intricate or compelling a task is, it will fail to provide a useful signal for benchmarking if it is too easy, too hard, unsolvable from the given dataset, or does not clearly distinguish between more or less effective methods. In this work, we carefully curate and adjust the difficulty of each task and dataset such that they can provide effective guidance for algorithms research. For some tasks, we provide multiple versions with varying difficulty, all the way up to tasks that are difficult to solve with current methods, so that the same benchmark can continuously be used to develop new methods even in the future. Our rule of thumb is to have, for each type of tasks, at least one task where the current state-of-the-art offline GCRL method achieves a success rate of 20-30%. This ensures that the task is solvable from the dataset while leaving significant room for improvement.

(3) Controllable datasets: The benchmark should provide tools to control and adjust datasets for scientific research and ablation studies. Verifying the effectiveness of algorithms in real-world

problems is surely important. However, for algorithms research, it is equally, if not more, important to provide analysis tools to enable a rigorous, scientific understanding of challenges and algorithms. Hence, instead of employing fixed, human-collected data, which does not always provide clear benchmarking signals for algorithms research, we choose to focus on simulated environments and synthetic datasets that we have full control of, and provide tools to reproduce and adjust them easily. We note that many algorithmic ideas in RL that have made a major impact, such as DQN (Mnih et al., 2013), PPO (Schulman et al., 2017), and CQL (Kumar et al., 2020), were originally developed in simulated environments. Even in natural language processing, studies on synthetic, controlled datasets have revealed the mechanisms and limitations of language models with scientific evidence (Allen-Zhu & Li, 2023a), and these insights have transferred to real scenarios (Allen-Zhu & Li, 2023b). We demonstrate how such controllability of datasets reveals challenges and design principles in offline GCRL in Section 8.2.

(4) Minimal unnecessary compute requirements: The tasks should be designed to minimize *unnecessary* computational overhead so that as many researchers as possible, including those from small labs and underprivileged backgrounds, can quickly iterate on their new *algorithmic* ideas. This does not mean that the tasks should be easy (indeed, some of our tasks are very challenging (but solvable)!); it means the benchmark should focus mainly on algorithmic challenges (*e.g.*, not requiring high-resolution image processing). In our benchmark, we provide both state- and pixel-based observations whenever possible, and minimize the size of image observations (up to $64 \times 64 \times 3$) to reduce the computational burden. Moreover, we carefully adjust colors, transparency, and lighting for image-based tasks to enable pixel-based control without high-resolution images or multiple views.

(5) High code quality: The reference implementations should be very clean and well-tuned so that researchers can directly use our implementations to build their ideas, and the benchmark should be very easy to set up. For reference implementations, we minimize the number of file dependencies for each algorithm, largely following the spirit of the single-file implementations of recent RL libraries (Huang et al., 2022; Tarasov et al., 2023), while maintaining a minimal amount of additional modularity. We also extensively test and tune different design choices and hyperparameters for each offline GCRL algorithm, to the degree that several methods achieve even better performances than their original performances reported on previous benchmarks (Table 3).

7 ENVIRONMENTS, TASKS, AND DATASETS


We now introduce the environments, tasks, and datasets in our benchmark. They can be broadly categorized into three groups: locomotion, manipulation, and drawing. We provide a separate validation dataset for each dataset, and most tasks support *both* state- and pixel-based observations. Videos are available at <https://ogbenchauthors.github.io/ogbench-anon/>.

Evaluation. Each task in OGBench accompanies five pre-defined state-goal pairs for evaluation (Appendix F). Performance is measured by the average success rate across the five evaluation goals. For each pre-define state-goal pair, we perform multiple rollouts with slightly randomized initial and goal states. In each evaluation episode, a goal $g \in \mathcal{S}$ (which is simply another state; see Section 2) is given to the agent, and the episode immediately terminates when the agent reaches the goal. Each task has its own goal success criterion, which we describe in Appendix E.1.


7.1 LOCOMOTION TASKS

We provide three types of locomotion environments, **AntMaze**, **HumanoidMaze**, and **AntSoccer**, with diverse variants. These environments are based on the MuJoCo simulator (Todorov et al., 2012).

AntMaze (antmaze) and **HumanoidMaze** (humanoidmaze). Maze navigation is one of the most widely used tasks for benchmarking offline GCRL algorithms. We provide two different types of maze navigation tasks, AntMaze, which involves controlling a quadrupedal Ant agent with 8 degrees of freedom (DoF), and HumanoidMaze, which involves controlling a much more complex 21-DoF Humanoid agent. The aim of these tasks is to learn to control a locomotion agent to reach a goal location in a given maze, based on potentially suboptimal navigation trajectories. Hence, the agent must learn *both* the maze structure and locomotion skills that involve high-dimensional control, purely from diverse offline trajectories.



In our benchmark, we substantially extend the original AntMaze tasks proposed by D4RL (Fu et al., 2020). Unlike the original D4RL tasks, which only support AntMaze and do not challenge stitching⁴, stochasticity, or pixel-based control, we support Humanoid control, pixel-based observations, and multi-goal evaluation, while providing more challenging and diverse types of mazes and datasets. The supported maze types are as follows:



- **medium:** This is the smallest maze, with the same layout as the original medium maze in D4RL.
- **large:** This is a larger maze, with the same layout as the original large maze in D4RL.
- **giant:** This is the largest maze, twice the size of large. It has the same size as the previous antmaze-ultra maze by Jiang et al. (2023), but its layout is more challenging and contains longer paths that require up to 3000 environment steps (for Humanoid). This maze is designed to substantially challenge the long-horizon reasoning capability of the agent.
- **teleport:** This maze is designed to challenge the agent’s ability to handle environment stochasticity. It has the same size as large, but contains multiple *stochastic* teleporters. If the agent enters a black hole, it is immediately sent to a randomly chosen white hole. However, since one of the three white holes is a dead-end, there is always a risk in taking a teleporter. The agent must therefore correctly handle stochasticity, without being optimistically biased.


On these mazes, we collect datasets with a low-level directional policy trained via SAC (Haarnoja et al., 2018) and a high-level waypoint controller. For each maze type, we provide three types of datasets that present different kinds of challenges: The figure on the right shows some example trajectories in these datasets.



- **navigate:** This is the regular dataset, collected by a noisy expert policy that navigates the maze by repeatedly reaching randomly sampled goals.
- **stitch:** This dataset is designed to challenge the agent’s stitching ability. It consists of short goal-reaching trajectories, where the length of each trajectory is at most 4 cell units. Hence, the agent must be able to stitch multiple trajectories (up to 8) together to complete the tasks.
- **explore:** This dataset is designed to test whether the agent can learn navigation skills from extremely low-quality (yet high-coverage) data. It consists of random exploratory trajectories, collected by commanding the low-level policy with random directions re-sampled every 10 steps, with a large amount of action noise.


We provide two types of observation modalities:

- **States:** This is the default setting, where the agent has access to the full low-dimensional state representation, including its current x - y position.
- **Pixels (visual):** This requires pure pixel-based control, where the agent only receives $64 \times 64 \times 3$ RGB images rendered from a third-person camera viewpoint. Following Park et al. (2023), we color the floor to enable the agent to infer its location from the images, obviating the need for a potentially expensive memory component.




⁴See Ghugare et al. (2024) for the details about this point.

AntSoccer (antsoccer). To provide a more diverse type of locomotion task beyond simple maze navigation, we introduce a new locomotion task, AntSoccer. This task involves controlling an Ant agent to dribble a soccer ball to a designated location in a maze. It is inspired by the quadruped-fetch task in the DeepMind Control suite (Tassa et al., 2018). AntSoccer is significantly harder than AntMaze because the agent must also carefully control the ball while navigating the environment. We provide two maze types: arena, which is an open space without walls, and medium, which is the same maze as the medium one in AntMaze. For datasets, we provide *navigate* and *stitch*. The *navigate* datasets consist of trajectories where the agent repeatedly approaches the ball and dribbles it to random locations. The *stitch* datasets consist of two different types of trajectories, maze navigation without the ball and dribbling with the ball near the agent, so that stitching is required to complete the full task. AntSoccer only supports state-based observations.




7.2 MANIPULATION TASKS


We provide a manipulation suite with three types of robotic manipulation tasks, **Cube**, **Scene**, and **Puzzle**, with diverse difficulties and complexities. They are respectively designed to test the agent’s regular object manipulation, sequential generalization, and combinatorial generalization abilities. These environments are based on the MuJoCo simulator (Todorov et al., 2012) and a 6-DoF UR5e robot arm (Zakka et al., 2022). For all manipulation tasks, we support *both* state-based observations and pixel-based observations with $64 \times 64 \times 3$ RGB camera images. For pixel images, we adjust colors and make the arm transparent to ensure full observability. The transparent arm in the figure might *appear* challenging, but the colors and transparency are carefully adjusted to minimize difficulties in visual perception, to the extent that some methods achieve even better performance with pixels (see Table 2).




Cube (cube). This task involves pick-and-place manipulation of cube blocks, whose goal is to control a robot arm to arrange cubes into designated configurations. We provide four variants, *single*, *double*, *triple*, and *quadruple*, with different numbers (1-4) of cubes. We provide “play”-style datasets (Lynch et al., 2019) collected by a scripted policy that repeatedly picks a random block and places it in other random locations or on another block. At test time, the agent is given goal configurations that require moving, stacking, swapping, or permuting cube blocks. Hence, the agent must learn not only generalizable multi-object pick-and-place behaviors from unstructured random trajectories in the dataset, but also long-term plans to achieve the tasks (*e.g.*, permuting blocks requires non-trivial sequential and logical reasoning).



Scene (scene). This task is designed to challenge the sequential, long-horizon reasoning capabilities of the agent. It involves manipulating diverse everyday objects, such as a cube block, a window, a drawer, and two button locks, where pressing a button toggles the lock status of the corresponding object (the drawer or window). We provide a “play”-style dataset collected by a scripted policy that randomly interacts with these objects. At test time, the agent is commanded to arrange the objects into a desired configuration. Evaluation tasks require a significant degree of sequential reasoning: for instance, some tasks require unlocking the drawer, opening it, putting the cube in the drawer, and closing it again (see the figure on the right), and the longest task involves seven atomic behaviors. Hence, the agent must be able to plan and sequentially combine learned manipulation skills.




Puzzle (puzzle). This task is designed to test the combinatorial generalization abilities of the agent. It requires solving the “Lights Out” puzzle⁵ with a robot arm. The puzzle consists of a two-dimensional array of buttons (e.g., a 4×6 grid), where pressing a button toggles the colors of the pressed button and the buttons adjacent to it (typically four, except on the edges and corners). The goal is to achieve a desired configuration of colors (e.g., turning all the buttons blue) by pressing an appropriate combination of buttons. Since these buttons are implemented in the MuJoCo simulator, the agent must control a robot arm to *physically* press the buttons. We provide four levels of difficulty, 3×3 , 4×4 , 4×5 , and 4×6 , with different grid sizes. The datasets are collected by a scripted policy that randomly presses buttons in arbitrary sequences. Given the enormous state space of this task (with up to $2^{24} = 16,777,216$ distinct button states), the agent must achieve *combinatorial* generalization while mastering low-level continuous control. Some evaluation task in the hardest puzzle requires pressing more than 20 buttons, which also substantially challenges the long-horizon reasoning capabilities of the agent. This might sound very challenging (and it is!), but we provide different levels and enough data to ensure that they provide meaningful research signals and are solvable (see the results in Section 8).



puzzle

7.3 DRAWING TASKS

Powderworld (powderworld). To provide more diverse tasks beyond robotic locomotion or manipulation, we introduce a *drawing* task, Powderworld (Frans & Isola, 2023)⁶, which presents unique challenges with extremely high intrinsic dimensionality. The goal of Powderworld is to draw a target picture on a 32×32 grid using different types of “powder” brushes, where each powder brush has a distinct physical property corresponding to a unique element. For example, the “sand” brush falls down and piles up, and the “fire” brush burns combustible elements like “plant.” We provide three versions of tasks, easy, medium, and hard, with different numbers of available elements (2, 5, and 8 elements, respectively). The datasets are collected by a random policy that keeps drawing arbitrary shapes with random brushes. This Powderworld task poses unique challenges that are distinct from the other tasks in the benchmark. First, the agent must deal with the high *intrinsic* dimensionality of the states, which presents a substantial challenge in representation learning. Second, since the transitions of powder elements are mostly stochastic and unpredictable, the agent must be able to correctly handle environment stochasticity. Third, the agent must achieve a high degree of generalization and sequential reasoning through a deep understanding of the physics, in order to complete symmetrical, orderly test-time drawing tasks from random, chaotic data.



powderworld

8 RESULTS

We now present and discuss the benchmarking results of existing offline goal-conditioned RL algorithms on OGBench.

8.1 ALGORITHMS

We benchmark six representative offline GCRL algorithms: goal-conditioned behavioral cloning (GCBC) (Lynch et al., 2019; Ghosh et al., 2021), goal-conditioned implicit $\{V, Q\}$ -learning (GCIVL and GCIQL) (Kostrikov et al., 2022; Park et al., 2023), quasimetric RL (QRL) (Wang et al., 2023), contrastive RL (CRL) (Eysenbach et al., 2022), and hierarchical implicit Q-learning (HIQL) (Park et al., 2023). GCBC is the simplest goal-conditioned behavioral cloning method. GCIVL and GCIQL are offline GCRL algorithms that approximate the optimal value function using an expectile regression (Newey & Powell, 1987). QRL is a non-traditional GCRL method that fits a quasimetric value function with a dual objective. CRL is a “one-step” RL algorithm that fits a Monte Carlo value function via contrastive learning and performs one-step policy improvement. HIQL is a hierarchical RL method that extracts a two-level hierarchical policy from a single GCIVL value function. For benchmarking, we perform a similar amount of hyperparameter tuning for each method to ensure fair comparison. We refer the reader to Appendix E for the details.

⁵[https://en.wikipedia.org/wiki/Lights_Out_\(game\)](https://en.wikipedia.org/wiki/Lights_Out_(game))

⁶Play here: <https://kvfrans.com/powder/>

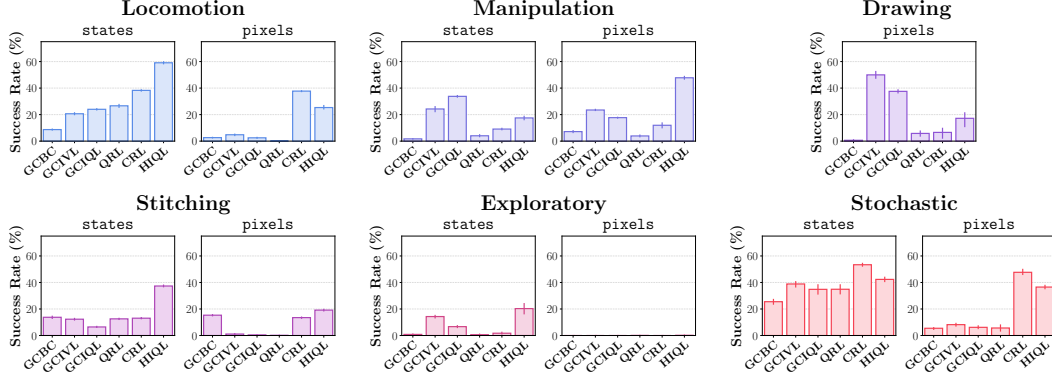


Figure 2: Benchmarking offline GCRL methods. We report the performances of six offline GCRL methods (GCBC, GCIVL, GCIQL, QRL, CRL, and HIQL), aggregated by different dataset categories (see Table 5 for the category list). The results are averaged over the tasks in each category, and then over 8 seeds (4 seeds for pixel-based tasks). Error bars denote 95% bootstrap confidence intervals. See Table 2 for the full results. In general, HIQL (a method that involves hierarchical policy extraction) tends to achieve strong performance across the board. Among the non-hierarchical methods, CRL tends to work best in locomotion tasks and GCIVL and GCIQL tend to work best in the others.

8.2 BENCHMARKING RESULTS AND Q&As

We present the full benchmarking results in Table 2 in Appendix B. Figure 2 summarizes the results by showing performances grouped by different task categories. Performances are measured by average (binary) success rates on the five test-time goals of each task. We train the agents for 1M gradient steps (500K for pixel-based tasks), and average the results over 8 seeds (4 seeds for pixel-based tasks). We discuss the results through Q&As (we refer to Appendix B for **more Q&As**).

Q: Which method works best in general?

A: While no single method dominates the others across all categories in Figure 2, HIQL (a method that involves hierarchical policy extraction) tends to achieve particularly strong performance among the benchmarked methods, especially in locomotion and visual manipulation tasks. Among the non-hierarchical methods, CRL tends to work best in locomotion tasks and GCIQL tends to work best in manipulation tasks. In the drawing tasks, GCIVL performs the best.

Q: There seem to be a lot of datasets. What should I use for my research?

A: For general offline GCRL algorithms research, we recommend starting with more “regular” tasks such as `antmaze-{large, giant}-navigate`, `cube-{single, double}`, and `scene`. From there, depending on the performance on these tasks, try harder versions of them or more challenging tasks like `humanoidmaze`, `antsoccer`, `puzzle`, and `powderworld`.

We also provide more specialized datasets that pose specific challenges in offline GCRL (Section 5). To test stitching, consider the `stitch` datasets in the locomotion suite as well as complex manipulation tasks that require stitching (e.g., `puzzle`). For long-horizon reasoning, `humanoidmaze-giant` has the longest episode length and `puzzle-4x6` has the most semantic steps. For stochasticity, try `antmaze-teleport`, which is specifically designed to challenge optimistically biased methods, and `powderworld`, which has unpredictable, stochastic dynamics. For highly suboptimal datasets, consider `antmaze-explore`, which features high suboptimality and high coverage.

Q: Have you found any insights on data collection for offline GCRL?

A: One of the main features of OGBench is that every task is accompanied by a reproducible and controllable data-generation script. Here, we show one example of how this controllability can lead to practical insights and raise open research questions. In Figure 3, we ablate the strength of Gaussian action noise σ added to expert actions on two manipulation tasks (`cube-single` and `puzzle-3x3`), and measure how this affects performance.




Figure 3: Datasets must be noisy enough.

The results are quite remarkable: they show that having the right amount of noise (i.e., state coverage) is *extremely* important for achieving good performance. For example, the performance drops from

99% to 6% if there is no noise in expert actions, even on the most basic cube pick-and-place task. This suggests that, we may need to prioritize coverage much more than optimality when collecting datasets for offline GCRL in the real world as well, and failing to do so may lead to (surprising) failures in learning. Like action noise, we believe there are many other important properties of datasets that significantly affect performance. We believe that our fully transparent, controllable data-generation scripts can facilitate such scientific studies.

9 RESEARCH OPPORTUNITIES

In this section, we discuss potential research ideas and open questions.

Be the first to solve unsolved tasks! While all environments in OGBench have at least one variant that current methods can solve to some degree, there are still a number of challenging tasks on which no existing method achieves non-trivial performance, such as `humanoidmaze-giant`, `cube-triple`, `puzzle-4x5`, `powderworld-hard`, and more. We ensure that sufficient data is available for those tasks (which is estimated from the amount needed to solve their easier versions). We invite researchers to take on these challenges and push the limits of offline GCRL with better algorithms.

How can we develop a policy that *generalizes well at test time*? In our experiments, we found hierarchical RL methods (*e.g.*, HIQL) to work especially well in several tasks. Among several potential explanations, we hypothesize that this is mainly because hierarchical RL reduces learning complexity by having two policies specialized in different things, which makes both policies *generalize* better at evaluation time. After all, test-time generalization is known to be one of the major bottlenecks in offline RL (Park et al., 2024a). But, are hierarchies really necessary to achieve good test-time generalization? Can we develop a non-hierarchical method that enjoys the same benefit by exploiting the subgoal structure of offline GCRL? This would be especially beneficial, not just because it is simpler, but also because it can potentially yield better, unified *representations* that can potentially serve as a “foundation model” for fine-tuning.

Can we develop a method that works well across all categories? Our benchmarking results reveal that no method consistently performs best across the board. HIQL tends to achieve strong performance but struggles in pixel-based locomotion and state-based manipulation. GCIQL shows strong performance in state-based manipulation, but struggles in locomotion. CRL exhibits the opposite trend: it excels in locomotion but underperforms in manipulation. Is there a way to combine only the strengths of these methods to develop a single approach that achieves the best performance across all types of tasks?

10 OUTLOOK

In this work, we introduced OGBench, a new benchmark designed to advance algorithms research in offline goal-conditioned RL. With the experimental results, we now revisit the very first question posed in this paper: “Why offline goal-conditioned RL?”

We hypothesize that offline goal-conditioned holds significant potential as a recipe for *general-purpose RL pre-training*, which yields richer and more diverse behaviors than (arguably more prevalent) generative pre-training, such as behavioral cloning and next-token prediction. Our experiments, albeit preliminary, show that even current offline GCRL algorithms can to some extent acquire effective policies for exceptionally long-horizon tasks entirely with sparse rewards, using data that is highly suboptimal. These results are not limited to toy example domains, but show up across a range of realistic simulated robotics and game-like settings in our benchmark. While generative objectives might capture the data distribution, offline GCRL can learn policies that actually achieve complex outcomes (such as beating a puzzle game) that could not be achieved simply by copying random data.

However, current offline GCRL algorithms also have various limitations. As shown in our results, they often struggle with long-horizon, high-dimensional tasks and those that require stitching, and no single method consistently outperforms others across all tasks. This suggests that we have not yet found the ideal algorithm that can fully realize the promise of offline GCRL as general-purpose RL pre-training. The first step toward finding such an ideal algorithm is to set up a solid benchmark that sufficiently challenges the limits of offline GCRL from diverse perspectives. We believe OGBench provides this foundation, and will lead to the development of performant, scalable offline GCRL algorithms that enable building foundation models for general-purpose behaviors.

REPRODUCIBILITY STATEMENT

We provide the full implementation details in Appendix E and the full code in the supplementary material.

REFERENCES

- Zeyuan Allen-Zhu and Yuanzhi Li. Physics of language models: Part 1, context-free grammar. *ArXiv*, abs/2305.13673, 2023a.
- Zeyuan Allen-Zhu and Yuanzhi Li. Physics of language models: Part 3.2, knowledge manipulation. *ArXiv*, abs/2309.14402, 2023b.
- Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. In *Neural Information Processing Systems (NeurIPS)*, 2017.
- Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. *ArXiv*, abs/1607.06450, 2016.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- G. Brockman, Vicki Cheung, Ludwig Pettersson, J. Schneider, John Schulman, Jie Tang, and W. Zaremba. OpenAI Gym. *ArXiv*, abs/1606.01540, 2016.
- Elliot Chane-Sane, Cordelia Schmid, and Ivan Laptev. Goal-conditioned reinforcement learning with imagined subgoals. In *International Conference on Machine Learning (ICML)*, 2021.
- Peter Dayan and Geoffrey E. Hinton. Feudal reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 1992.
- Lasse Espeholt, Hubert Soyer, Rémi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *International Conference on Machine Learning (ICML)*, 2018.
- Benjamin Eysenbach, Ruslan Salakhutdinov, and Sergey Levine. Search on the replay buffer: Bridging planning and reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2019.
- Benjamin Eysenbach, Tianjun Zhang, Ruslan Salakhutdinov, and Sergey Levine. Contrastive learning as goal-conditioned reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2022.
- Kuan Fang, Patrick Yin, Ashvin Nair, and Sergey Levine. Planning to practice: Efficient online fine-tuning by composing goals in latent space. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- Kevin Frans and Phillip Isola. Powderworld: A platform for understanding generalization via rich task distributions. In *International Conference on Learning Representations (ICLR)*, 2023.
- Justin Fu, Aviral Kumar, Ofir Nachum, G. Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning. *ArXiv*, abs/2004.07219, 2020.
- Scott Fujimoto and Shixiang Shane Gu. A minimalist approach to offline reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- Dibya Ghosh, Abhishek Gupta, Ashwin Reddy, Justin Fu, Coline Devin, Benjamin Eysenbach, and Sergey Levine. Learning to reach goals via iterated supervised learning. In *International Conference on Learning Representations (ICLR)*, 2021.

- Dibya Ghosh, Chethan Bhateja, and Sergey Levine. Reinforcement learning from passive data via latent intentions. In *International Conference on Machine Learning (ICML)*, 2023.
- Raj Ghugare, Matthieu Geist, Glen Berseth, and Benjamin Eysenbach. Closing the gap between td learning and supervised learning—a generalisation point of view. In *International Conference on Learning Representations (ICLR)*, 2024.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning (ICML)*, 2018.
- Joey Hejna, Jensen Gao, and Dorsa Sadigh. Distance weighted supervised learning for offline interaction data. In *International Conference on Machine Learning (ICML)*, 2023.
- Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *ArXiv*, abs/1606.08415, 2016.
- Christopher Hoang, Sungryull Sohn, Jongwook Choi, Wilka Carvalho, and Honglak Lee. Successor feature landmarks for long-horizon goal-conditioned reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- Mineui Hong, Minjae Kang, and Songhwai Oh. Diffused task-agnostic milestone planner. In *Neural Information Processing Systems (NeurIPS)*, 2023.
- Shengyi Huang, Rousslan Fernand Julien Dossa, Chang Ye, Jeff Braga, Dipam Chakraborty, Kinal Mehta, and JoÅGo GM AraÅšjo. Cleanrl: High-quality single-file implementations of deep reinforcement learning algorithms. *Journal of Machine Learning Research (JMLR)*, 23(274):1–18, 2022.
- Zhiao Huang, Fangchen Liu, and Hao Su. Mapping state space using landmarks for universal goal reaching. In *Neural Information Processing Systems (NeurIPS)*, 2019.
- Zhengyao Jiang, Tianjun Zhang, Michael Janner, Yueying Li, Tim Rocktaschel, Edward Grefenstette, and Yuandong Tian. Efficient planning in a compact latent action space. In *International Conference on Learning Representations (ICLR)*, 2023.
- Leslie Pack Kaelbling. Learning to achieve goals. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 1993.
- Junsu Kim, Younggyo Seo, and Jinwoo Shin. Landmark-guided subgoal generation in hierarchical reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- Junsu Kim, Seohong Park, and Sergey Levine. Unsupervised-to-online reinforcement learning. *ArXiv*, abs/2408.14785, 2024.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline reinforcement learning with implicit q-learning. In *International Conference on Learning Representations (ICLR)*, 2022.
- Aviral Kumar, Aurick Zhou, G. Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2020.
- Jinning Li, Chen Tang, Masayoshi Tomizuka, and Wei Zhan. Hierarchical planning through goal-conditioned offline reinforcement learning. *IEEE Robotics and Automation Letters (RA-L)*, 7(4): 10216–10223, 2022.
- Bo Liu, Yihao Feng, Qiang Liu, and Peter Stone. Metric residual network for sample efficient goal-conditioned reinforcement learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2023.
- Corey Lynch, Mohi Khansari, Ted Xiao, Vikash Kumar, Jonathan Tompson, Sergey Levine, and Pierre Sermanet. Learning latent plans from play. In *Conference on Robot Learning (CoRL)*, 2019.

- Yecheng Jason Ma, Jason Yan, Dinesh Jayaraman, and Osbert Bastani. How far i'll go: Offline goal-conditioned reinforcement learning via f-advantage regression. In *Neural Information Processing Systems (NeurIPS)*, 2022.
- Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy Zhang. Vip: Towards universal visual reward and representation via value-implicit pre-training. In *International Conference on Learning Representations (ICLR)*, 2023.
- Zhuang Ma and Michael Collins. Noise contrastive estimation and negative sampling for conditional models: Consistency and statistical efficiency. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin A. Riedmiller. Playing atari with deep reinforcement learning. *ArXiv*, abs/1312.5602, 2013.
- Vivek Myers, Chongyi Zheng, Anca Dragan, Sergey Levine, and Benjamin Eysenbach. Learning temporal distances: Contrastive successor features can provide a metric structure for decision-making. In *International Conference on Machine Learning (ICML)*, 2024.
- Soroush Nasiriany, Vitchyr H. Pong, Steven Lin, and Sergey Levine. Planning with goal-conditioned policies. In *Neural Information Processing Systems (NeurIPS)*, 2019.
- Whitney Newey and James L. Powell. Asymmetric least squares estimation and testing. *Econometrica*, 55:819–847, 1987.
- Seohong Park, Dibya Ghosh, Benjamin Eysenbach, and Sergey Levine. Hiql: Offline goal-conditioned rl with latent states as actions. In *Neural Information Processing Systems (NeurIPS)*, 2023.
- Seohong Park, Kevin Frans, Sergey Levine, and Aviral Kumar. Is value learning really the main bottleneck in offline rl? In *Neural Information Processing Systems (NeurIPS)*, 2024a.
- Seohong Park, Tobias Kreiman, and Sergey Levine. Foundation policies with hilbert representations. In *International Conference on Machine Learning (ICML)*, 2024b.
- Seohong Park, Oleh Rybkin, and Sergey Levine. Metra: Scalable unsupervised rl with metric-aware abstraction. In *International Conference on Learning Representations (ICLR)*, 2024c.
- Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *ArXiv*, abs/1910.00177, 2019.
- Jan Peters and Stefan Schaal. Reinforcement learning by reward-weighted regression for operational space control. In *International Conference on Machine Learning (ICML)*, 2007.
- Matthias Plappert, Marcin Andrychowicz, Alex Ray, Bob McGrew, Bowen Baker, Glenn Powell, Jonas Schneider, Josh Tobin, Maciek Chociej, Peter Welinder, et al. Multi-goal reinforcement learning: Challenging robotics environments and request for research. *ArXiv*, abs/1802.09464, 2018.
- Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun. Semi-parametric topological memory for navigation. In *International Conference on Learning Representations (ICLR)*, 2018.
- Tom Schaul, Dan Horgan, Karol Gregor, and David Silver. Universal value function approximators. In *International Conference on Machine Learning (ICML)*, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *ArXiv*, abs/1707.06347, 2017.
- Harshit Sikchi, Rohan Chitnis, Ahmed Touati, Alborz Geramifard, Amy Zhang, and Scott Niekum. Score models for offline goal-conditioned reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2024.
- Denis Tarasov, Alexander Nikulin, Dmitry Akimov, Vladislav Kurenkov, and Sergey Kolesnikov. Corl: Research-oriented deep offline reinforcement learning library. In *Neural Information Processing Systems (NeurIPS)*, 2023.

- Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy P. Lillicrap, and Martin A. Riedmiller. Deepmind control suite. *ArXiv*, abs/1801.00690, 2018.
- Jordan Terry, Benjamin Black, Nathaniel Grammel, Mario Jayakumar, Ananth Hari, Ryan Sullivan, Luis S Santos, Clemens Dieffendahl, Caroline Horsch, Rodrigo Perez-Vicente, et al. Pettingzoo: Gym for multi-agent reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- Mark Towers, Ariel Kwiatkowski, Jordan Terry, John U Balis, Gianluca De Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Markus Krimmel, Arjun KG, et al. Gymnasium: A standard interface for reinforcement learning environments. *ArXiv*, abs/2407.17032, 2024.
- Mianchu Wang, Rui Yang, Xi Chen, and Meng Fang. Goplan: Goal-conditioned offline reinforcement learning by planning with learned models. *Transactions on Machine Learning Research (TMLR)*, 2024.
- Tongzhou Wang and Phillip Isola. Improved representation of asymmetrical distances with interval quasimetric embeddings. *ArXiv*, abs/2211.15120, 2022.
- Tongzhou Wang, Antonio Torralba, Phillip Isola, and Amy Zhang. Optimal goal-reaching reinforcement learning via quasimetric learning. In *International Conference on Machine Learning (ICML)*, 2023.
- Rui Yang, Yiming Lu, Wenzhe Li, Hao Sun, Meng Fang, Yali Du, Xiu Li, Lei Han, and Chongjie Zhang. Rethinking goal-conditioned supervised learning and its connection to offline rl. In *International Conference on Learning Representations (ICLR)*, 2022.
- Rui Yang, Yong Lin, Xiaoteng Ma, Haotian Hu, Chongjie Zhang, and T. Zhang. What is essential for unseen goal generalization of offline goal-conditioned rl? In *International Conference on Machine Learning (ICML)*, 2023.
- Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning (CoRL)*, 2019.
- Kevin Zakka, Yuval Tassa, and MuJoCo Menagerie Contributors. Mujoco menagerie: A collection of high-quality simulation models for mujoco, 2022. URL http://github.com/google-deepmind/mujoco_menagerie.
- Zilai Zeng, Ce Zhang, Shijie Wang, and Chen Sun. Goal-conditioned predictive coding for offline reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2023.
- Chongyi Zheng, Benjamin Eysenbach, Homer Walke, Patrick Yin, Kuan Fang, Ruslan Salakhutdinov, and Sergey Levine. Stabilizing contrastive rl: Techniques for offline goal reaching. In *International Conference on Learning Representations (ICLR)*, 2024a.
- Chongyi Zheng, Ruslan Salakhutdinov, and Benjamin Eysenbach. Contrastive difference predictive coding. In *International Conference on Learning Representations (ICLR)*, 2024b.

Table 2: **Full benchmarking results.** We report each method’s average (binary) success rate (%) and standard deviation across the five test-time goals on each task. The results are averaged over 8 seeds (4 seeds for pixel-based tasks).

Environment	Dataset Type	Dataset	GCBC	GCIVL	GCIQL	QRL	CRL	HIQL
antmaze	navigate	antmaze-medium-navigate-v0	29.0 \pm 4.0	72.0 \pm 7.9	71.2 \pm 4.1	87.9 \pm 3.5	94.8 \pm 1.3	95.9 \pm 0.9
		antmaze-large-navigate-v0	24.0 \pm 1.8	15.7 \pm 5.4	34.2 \pm 3.6	74.6 \pm 6.4	82.8 \pm 4.0	91.0 \pm 2.2
		antmaze-giant-navigate-v0	0.2 \pm 0.2	0.2 \pm 0.2	0.2 \pm 0.1	13.7 \pm 3.4	15.7 \pm 3.3	65.2 \pm 5.3
		antmaze-teleport-navigate-v0	25.5 \pm 2.8	38.9 \pm 3.0	34.8 \pm 5.4	34.9 \pm 5.3	53.4 \pm 1.6	42.3 \pm 2.5
	stitch	antmaze-medium-stitch-v0	45.5 \pm 11.0	44.0 \pm 5.6	29.3 \pm 6.2	59.1 \pm 6.7	52.6 \pm 6.2	93.8 \pm 0.9
		antmaze-large-stitch-v0	3.4 \pm 2.8	18.5 \pm 2.4	7.5 \pm 2.0	18.4 \pm 2.1	10.8 \pm 1.8	66.7 \pm 5.3
		antmaze-giant-stitch-v0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.4 \pm 0.5	0.0 \pm 0.0	1.8 \pm 1.6
		antmaze-teleport-stitch-v0	31.3 \pm 6.3	39.1 \pm 3.0	16.7 \pm 2.3	23.7 \pm 5.2	30.5 \pm 3.6	35.9 \pm 2.5
	explore	antmaze-medium-explore-v0	1.6 \pm 1.0	18.8 \pm 2.7	13.0 \pm 2.1	1.1 \pm 1.3	3.1 \pm 2.1	36.7 \pm 10.5
		antmaze-large-explore-v0	0.0 \pm 0.1	9.8 \pm 3.4	0.4 \pm 0.3	0.0 \pm 0.0	0.3 \pm 0.2	3.9 \pm 5.0
		antmaze-teleport-explore-v0	2.4 \pm 1.1	32.0 \pm 1.7	7.3 \pm 3.3	2.3 \pm 2.1	19.5 \pm 2.2	34.1 \pm 15.0
humanoidmaze	navigate	humanoidmaze-medium-navigate-v0	7.5 \pm 1.6	24.0 \pm 2.2	27.3 \pm 2.5	21.4 \pm 8.2	59.9 \pm 3.8	88.5 \pm 1.8
		humanoidmaze-large-navigate-v0	1.4 \pm 0.4	2.5 \pm 0.6	2.0 \pm 0.7	5.3 \pm 1.3	24.0 \pm 4.2	49.3 \pm 4.2
		humanoidmaze-giant-navigate-v0	0.2 \pm 0.1	0.2 \pm 0.3	0.5 \pm 0.3	0.7 \pm 0.3	3.1 \pm 1.5	12.5 \pm 4.2
	stitch	humanoidmaze-medium-stitch-v0	29.0 \pm 4.8	12.3 \pm 1.7	12.1 \pm 3.1	18.0 \pm 2.1	36.2 \pm 2.5	87.9 \pm 2.2
		humanoidmaze-large-stitch-v0	5.6 \pm 2.9	1.2 \pm 0.5	0.5 \pm 0.2	3.5 \pm 1.4	3.9 \pm 0.6	28.4 \pm 2.6
		humanoidmaze-giant-stitch-v0	0.1 \pm 0.1	0.2 \pm 0.1	0.2 \pm 0.1	0.4 \pm 0.3	0.2 \pm 0.1	3.3 \pm 1.9
antsoccer	navigate	antsoccer-arena-navigate-v0	5.2 \pm 0.6	46.6 \pm 2.6	49.9 \pm 2.4	7.6 \pm 1.8	22.6 \pm 2.2	57.6 \pm 2.4
		antsoccer-medium-navigate-v0	2.1 \pm 0.5	3.9 \pm 1.0	6.8 \pm 1.2	1.9 \pm 2.0	3.2 \pm 1.0	12.5 \pm 1.5
	stitch	antsoccer-arena-stitch-v0	24.5 \pm 7.8	21.4 \pm 3.1	2.1 \pm 0.4	0.8 \pm 0.5	0.7 \pm 0.3	15.3 \pm 1.3
visual-antmaze	navigate	visual-antmaze-medium-navigate-v0	11.4 \pm 1.5	22.4 \pm 2.4	10.8 \pm 1.4	0.0 \pm 0.0	94.3 \pm 0.6	92.7 \pm 3.8
		visual-antmaze-large-navigate-v0	3.9 \pm 0.5	5.5 \pm 1.0	3.9 \pm 0.8	0.3 \pm 0.4	83.9 \pm 1.2	52.6 \pm 8.6
		visual-antmaze-giant-navigate-v0	0.4 \pm 0.3	1.0 \pm 0.8	0.3 \pm 0.2	0.1 \pm 0.2	47.2 \pm 1.7	6.5 \pm 4.3
		visual-antmaze-teleport-navigate-v0	5.5 \pm 0.8	8.2 \pm 1.2	6.2 \pm 1.2	5.7 \pm 2.7	47.7 \pm 2.2	36.5 \pm 1.5
	stitch	visual-antmaze-medium-stitch-v0	67.3 \pm 3.6	5.9 \pm 2.1	2.4 \pm 0.3	0.1 \pm 0.2	69.0 \pm 2.2	86.7 \pm 2.4
		visual-antmaze-large-stitch-v0	23.6 \pm 2.7	0.8 \pm 0.6	0.1 \pm 0.1	0.6 \pm 0.6	11.1 \pm 2.5	27.9 \pm 2.1
		visual-antmaze-giant-stitch-v0	0.3 \pm 0.4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.1	0.2 \pm 0.2
		visual-antmaze-teleport-stitch-v0	31.8 \pm 3.1	1.4 \pm 0.8	1.0 \pm 0.4	1.4 \pm 1.6	31.7 \pm 6.4	37.0 \pm 3.5
	explore	visual-antmaze-medium-explore-v0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.1	0.0 \pm 0.0	0.2 \pm 0.4
		visual-antmaze-large-explore-v0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
		visual-antmaze-teleport-explore-v0	0.0 \pm 0.0	0.1 \pm 0.1	0.2 \pm 0.2	0.1 \pm 0.1	1.0 \pm 0.4	18.7 \pm 7.6
visual-humanoidmaze	navigate	visual-humanoidmaze-medium-navigate-v0	0.0 \pm 0.0	0.0 \pm 0.1	0.0 \pm 0.0	0.0 \pm 0.0	0.9 \pm 0.4	0.2 \pm 0.3
		visual-humanoidmaze-large-navigate-v0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
		visual-humanoidmaze-giant-navigate-v0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
	stitch	visual-humanoidmaze-medium-stitch-v0	0.9 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.8 \pm 0.4	0.2 \pm 0.3
		visual-humanoidmaze-large-stitch-v0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
		visual-humanoidmaze-giant-stitch-v0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
cube	play	cube-single-play-v0	8.7 \pm 1.6	70.3 \pm 5.3	99.4 \pm 0.4	24.7 \pm 8.1	40.4 \pm 4.6	42.0 \pm 7.5
		cube-double-play-v0	1.8 \pm 1.1	15.3 \pm 2.7	21.8 \pm 5.0	4.3 \pm 1.5	2.7 \pm 0.5	2.2 \pm 0.9
		cube-triple-play-v0	2.4 \pm 1.2	9.4 \pm 2.6	2.2 \pm 0.9	2.1 \pm 1.5	2.0 \pm 0.8	1.7 \pm 0.7
		cube-quadruple-play-v0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.1
scene	play	scene-play-v0	1.6 \pm 0.7	38.2 \pm 4.4	37.3 \pm 6.8	5.2 \pm 1.4	1.0 \pm 0.6	35.7 \pm 4.6
puzzle	play	puzzle-3x3-play-v0	0.9 \pm 0.5	47.1 \pm 20.8	95.7 \pm 1.1	0.0 \pm 0.1	35.8 \pm 6.3	57.4 \pm 11.1
		puzzle-4x4-play-v0	0.0 \pm 0.0	20.4 \pm 4.0	29.9 \pm 7.3	0.0 \pm 0.0	0.0 \pm 0.0	15.6 \pm 3.8
		puzzle-4x5-play-v0	0.0 \pm 0.0	0.7 \pm 0.6	0.3 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	1.0 \pm 0.4
		puzzle-4x6-play-v0	0.0 \pm 0.0	16.8 \pm 1.5	17.2 \pm 3.3	0.0 \pm 0.0	0.0 \pm 0.0	1.7 \pm 0.9
visual-cube	play	visual-cube-single-play-v0	17.3 \pm 4.2	74.9 \pm 1.7	42.1 \pm 2.9	11.3 \pm 9.9	59.5 \pm 16.1	98.4 \pm 0.5
		visual-cube-double-play-v0	3.7 \pm 1.2	13.6 \pm 1.5	8.6 \pm 1.0	9.5 \pm 4.4	3.8 \pm 1.4	58.3 \pm 2.7
		visual-cube-triple-play-v0	17.5 \pm 0.8	16.2 \pm 1.0	10.8 \pm 0.7	10.8 \pm 3.4	17.3 \pm 1.6	24.9 \pm 1.0
		visual-cube-quadruple-play-v0	8.0 \pm 1.0	0.2 \pm 0.3	0.3 \pm 0.2	1.8 \pm 2.5	7.2 \pm 1.4	17.0 \pm 1.6
visual-scene	play	visual-scene-play-v0	9.3 \pm 1.7	22.6 \pm 1.7	9.2 \pm 2.0	1.7 \pm 1.0	10.4 \pm 3.8	60.7 \pm 5.6
visual-puzzle	play	visual-puzzle-3x3-play-v0	0.1 \pm 0.1	20.1 \pm 0.1	22.8 \pm 1.3	0.0 \pm 0.0	0.6 \pm 1.1	69.0 \pm 12.4
		visual-puzzle-4x4-play-v0	5.4 \pm 1.9	47.1 \pm 2.1	47.9 \pm 4.5	0.0 \pm 0.0	6.8 \pm 1.7	87.5 \pm 5.0
		visual-puzzle-4x5-play-v0	0.4 \pm 0.4	0.0 \pm 0.0	0.3 \pm 0.2	0.0 \pm 0.0	0.4 \pm 0.2	1.4 \pm 0.5
		visual-puzzle-4x6-play-v0	2.9 \pm 1.7	16.6 \pm 1.6	17.5 \pm 0.9	0.0 \pm 0.0	1.5 \pm 1.0	12.8 \pm 5.1
powderworld	play	powderworld-easy-play-v0	0.2 \pm 0.3	89.7 \pm 4.7	97.7 \pm 1.3	15.1 \pm 6.5	18.4 \pm 12.8	24.7 \pm 18.0
		powderworld-medium-play-v0	1.6 \pm 0.9	56.6 \pm 6.6	14.9 \pm 2.9	2.4 \pm 1.2	1.3 \pm 0.9	25.8 \pm 2.4
		powderworld-hard-play-v0	0.1 \pm 0.1	3.5 \pm 2.3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	1.2 \pm 1.7

A LIMITATIONS

While OGBench covers a number of challenges in offline goal-conditioned RL, such as long-horizon reasoning, goal stitching, and stochastic control, there exist other challenges that our benchmark does not address. For example, all OGBench tasks assume that the environment dynamics remain the same between the training and evaluation environments. Also, although several OGBench tasks (*e.g.*, Cube, Puzzle, and Powderworld) require unseen goal generalization to some degree, our tasks do not specifically test visual generalization to entirely new objects. Finally, we have made several trade-offs to reduce computational cost and to focus the benchmark on algorithms research at the expense of sacrificing realism to some degree (*e.g.*, the use of the transparent arm in manipulation environments, the use of synthetic (yet fully controllable) datasets, etc.). Nonetheless, we believe OGBench can spur the development of performant offline GCRL *algorithms*, which can then help researchers develop scalable data-driven unsupervised RL pre-training methods for real-world tasks.

Table 3: **How good are our reference implementations?** Our implementations generally achieve better performance than previously reported ones.

D4RL antmaze-large-diverse-v2			D4RL antmaze-large-play-v2		
Method	Previously Reported Performance	Ours	Method	Previously Reported Performance	Ours
GCBC	20 (Park et al., 2023)	41 ± 7	GCBC	23 (Park et al., 2023)	39 ± 4
GCIVL	51 (Park et al., 2023)	64 ± 8	GCIVL	57 (Park et al., 2023)	58 ± 8
GCIQL	30 (Zeng et al., 2023)	64 ± 10	GCIQL	40 (Zeng et al., 2023)	55 ± 11
QRL	52 ⁷ (Zheng et al., 2024b)	37 ± 14	QRL	53 ⁷ (Zheng et al., 2024b)	38 ± 8
CRL	54 (Eysenbach et al., 2022)	79 ± 6	CRL	49 (Eysenbach et al., 2022)	74 ± 4
HIQL	88 (Park et al., 2023)	87 ± 3	HIQL	86 (Park et al., 2023)	87 ± 2

B ADDITIONAL RESULTS AND Q&AS

We present the full benchmarking results in Table 2, and provide additional Q&As in this section.

Q: Which methods are good at goal stitching?

A: To see this, we can compare the performances on the `navigate` and `stitch` datasets from the same locomotion task in Table 2. The results suggest that, as expected, full RL-based methods like HIQL (*i.e.*, methods that fit the optimal value function Q^*) are better at stitching than one-step RL methods like CRL (*i.e.*, methods that fit the behavioral value function Q^β). For example, in visual locomotion tasks, the relative performance between HIQL and CRL is reversed on the `stitch` datasets (Figure 2).

Q: Which methods are good at handling stochasticity?

A: For this, we can compare the performances on the `antmaze-large` and `antmaze-teleport` tasks in Table 2, where both have the same maze size, but only the latter involves stochastic transitions that incur risk. Table 2 shows that, value-only methods like HIQL and QRL (*i.e.*, methods that do not have a separate Q function), which are optimistically biased in stochastic environments, struggle relatively more in stochastic teleport tasks. In contrast, CRL is generally robust to environment stochasticity, likely because it fits a Monte Carlo value function.

Q: Which methods are good at handling pixel-based observations?

A: Although state-based and pixel-based observations generally provide the same amount of information, several methods struggle to handle image observations due to additional representational challenges. We can understand how well a method addresses such representational challenges by comparing the performances of corresponding state- and pixel-based tasks. Table 2 shows that CRL is notably robust to the difference in input modalities, likely because it is based on a pure representation learning objective. HIQL also achieves strong performance in pixel-based tasks, especially in visual manipulation tasks. However, these methods are still not perfect at handling image observations; for example, HIQL achieves relatively weak performance on image drawing tasks. We suspect this is due to the difficulty of learning low-dimensional subgoal representations from states of high intrinsic dimensionality.

Q: How good are our reference implementations?

A: To see this, we compare the performance of our reference implementations with previously reported numbers on one of the most commonly used tasks in prior work, D4RL `antmaze-large` (Fu et al., 2020). Table 3 shows the comparison results, with the corresponding numbers taken from the prior works (Eysenbach et al., 2022; Park et al., 2023; Zeng et al., 2023; Zheng et al., 2024b). In general, our implementations achieve better performance than previously reported results, sometimes significantly surpassing them (*e.g.*, CRL).

Q: Can I just use D4RL AntMaze instead of OGBench AntMaze?

A: D4RL AntMaze is an excellent task for benchmarking offline RL algorithms. However, it is limited for benchmarking offline *goal-conditioned* RL algorithms because it only involves a single, fixed state-goal pair, and the datasets are tailored to this specific task (Fu et al., 2020). In contrast, OGBench supports multi-goal evaluation (and provides much more diverse types of tasks and datasets!). To empirically demonstrate this difference, we compare the benchmarking results on

⁷We note that Zheng et al. (2024b) use a different evaluation scheme based on the best performance. We report the average performance over the last three evaluation epochs (see Appendix E).

Table 4: **Do not use single-goal evaluation!** Only using a single state-goal pair (a common practice in D4RL tasks) can potentially lead to inaccurate conclusions about offline GCRL methods. OGBench always uses multi-goal evaluation. See how the rank between GCIQL and QRL is reversed with multi-goal evaluation on the same antmaze-large maze.

Dataset	GCBC	GCIVL	GCIQL	QRL	CRL	HIQL
D4RL antmaze-large-diverse-v2 (single-goal evaluation)	41 \pm 7	64 \pm 8	64 \pm 10	37 \pm 14	79 \pm 6	87 \pm 3
D4RL antmaze-large-play-v2 (single-goal evaluation)	39 \pm 4	58 \pm 8	55 \pm 11	38 \pm 8	74 \pm 4	87 \pm 2
OGBench antmaze-large-navigate-v0 (multi-goal evaluation, ours)	24 \pm 2	16 \pm 5	34 \pm 4	75 \pm 6	83 \pm 4	91 \pm 2

D4RL antmaze-large- $\{\text{diverse}, \text{play}\}$ and OGBench antmaze-large-navigate in Table 4. The table suggests that single-goal evaluation is indeed limited, and is potentially prone to inaccurate conclusions: for example, see how the ranking between GCIQL and QRL is reversed with multi-goal evaluation on the same antmaze-large maze. Moreover, the performance differences between methods are more pronounced in OGBench AntMaze, showing that OGBench provides clearer research signals.

C PRIOR WORK IN GOAL-CONDITIONED RL

The problem of reaching any goal from any state has long been considered one of the central problems in reinforcement learning and sequential decision making (Kaelbling, 1993; Schaul et al., 2015; Andrychowicz et al., 2017), owing to its unsupervised nature, simplicity, and generality. There are many unique features of goal-conditioned RL that make it distinct from other (multi-task) RL problems, such as the presence of recursive subgoal structures, metric structures, and probabilistic interpretations. These intriguing properties have led to the development of diverse families of online and offline GCRL algorithms based on hindsight relabeling (Andrychowicz et al., 2017), hierarchical learning (Dayan & Hinton, 1992; Chane-Sane et al., 2021; Li et al., 2022; Park et al., 2023), planning (Savinov et al., 2018; Eysenbach et al., 2019; Nasiriany et al., 2019; Huang et al., 2019; Hoang et al., 2021; Kim et al., 2021; Wang et al., 2024), metric learning (Wang et al., 2023; Park et al., 2024b; Myers et al., 2024), dual optimization (Ma et al., 2022; 2023; Sikchi et al., 2024), weighted behavioral cloning (Yang et al., 2022; 2023; Hejna et al., 2023), generative modeling (Zeng et al., 2023; Hong et al., 2023), and contrastive learning (Eysenbach et al., 2022; Zheng et al., 2024a;b). In this work, we also consider the problem of offline goal-conditioned RL; however, instead of proposing a new algorithm, we introduce a new benchmark and reference implementations to facilitate and advance algorithms research in offline GCRL.

D OFFLINE GCRL ALGORITHMS

In this section, we describe the six offline GCRL methods used for benchmarking in detail. We first define four goal-sampling distributions that correspond to the current state, uniform future states, geometric future states, and random states, respectively:

- $p_{\text{cur}}^{\mathcal{D}}(g | s)$ denotes the Dirac delta distribution at s (i.e., g is always set to s).
- $p_{\text{traj}}^{\mathcal{D}}(g | s)$ denotes the uniform future state distribution defined as follows: assuming $s = s_t$ in a trajectory $\tau = (s_0, a_0, s_1, \dots, s_T)^8$, we sample an index k from the uniform distribution $\text{Unif}(\min(t+1, T-1), T-1)$ (inclusive), and set $g = s_k$.
- $p_{\text{geom}}^{\mathcal{D}}(g | s)$ denotes the truncated geometric future state distribution defined as follows: assuming $s = s_t$ in a trajectory $\tau = (s_0, a_0, s_1, \dots, s_T)$, we sample an index k from the geometric distribution $\text{Geom}(1 - \gamma)$ (whose support starts from 1), and set $g = s_{\min(s+k, T-1)}$.
- $p_{\text{rand}}^{\mathcal{D}}(g)$ denotes the uniform state distribution over the dataset \mathcal{D} .

Additionally, $p_{\text{mixed}}^{\mathcal{D}}(g | s)$ denotes a mixture of these four goal-sampling distributions with a mixture ratio defined by hyperparameters, $p^{\mathcal{D}}(\cdot)$ simply denotes the uniform distribution over the dataset, and we sometimes use $p^{\mathcal{D}}(\cdot | s, a)$ instead of $p^{\mathcal{D}}(\cdot | s)$ to denote the distribution corresponding to the state-action pair.

⁸If there are multiple such (τ, t) tuples in the dataset, consider the uniform mixture of them.

Goal-conditioned behavioral cloning (GCBC). GCBC (Lynch et al., 2019; Ghosh et al., 2021) simply performs behavioral cloning using future states in the same trajectory as goals. GCBC maximizes the following objective to train a goal-conditioned policy $\pi(a | s, g)$.

$$J_{\text{GCBC}}(\pi) = \mathbb{E}_{(s,a) \sim p^{\mathcal{D}}(s,a), g \sim p_{\text{traj}}^{\mathcal{D}}(g|s)} [\log \pi(a | s, g)]. \quad (1)$$

Goal-conditioned implicit {V, Q}-learning (GCIVL and GCIQL). GCIVL and GCIQL are goal-conditioned variants of implicit Q-learning (IQL) (Kostrikov et al., 2022), which is an offline RL algorithm that fits the optimal value functions (V^* or Q^*) using an expectile regression (Newey & Powell, 1987). GCIQL is a straightforward goal-conditioned variant of IQL, and GCIVL is the V -only variant introduced by Park et al. (2023). GCIVL fits a value function $V(s, g)$ by minimizing the following loss:

$$\mathcal{L}_{\text{GCIVL}}(V) = \mathbb{E}_{s \sim p^{\mathcal{D}}(s), g \sim p_{\text{mixed}}^{\mathcal{D}}(g|s)} [\ell_{\kappa}^2(r(s, g) + \gamma \bar{V}(s', g) - V(s, g))], \quad (2)$$

where $r(s, g) = \mathbf{1}_{\{g\}}(s) - 1$ denotes the $(-1, 0)$ -sparse goal-conditioned reward function, \bar{V} denotes the target value function (Mnih et al., 2013), and $\ell_{\kappa}^2(x) = |\kappa - \mathbf{1}_{\{x : x < 0\}}(x)|x^2$ denotes the expectile loss with an expectile κ .

GCIQL fits both $V(s, g)$ and $Q(s, a, g)$ by jointly minimizing the following losses:

$$\mathcal{L}_{\text{GCIQL}}^V(V) = \mathbb{E}_{(s,a) \sim p^{\mathcal{D}}(s,a), g \sim p_{\text{mixed}}^{\mathcal{D}}(g|s)} [\ell_{\kappa}^2(\bar{Q}(s, a, g) - V(s, g))], \quad (3)$$

$$\mathcal{L}_{\text{GCIQL}}^Q(Q) = \mathbb{E}_{(s,a,s') \sim p^{\mathcal{D}}(s,a,s'), g \sim p_{\text{mixed}}^{\mathcal{D}}(g|s)} [(r(s, g) + \gamma V(s', g) - Q(s, a, g))^2], \quad (4)$$

where \bar{Q} denotes the target Q function (Mnih et al., 2013). We note that GCIVL is optimistically biased in stochastic environments, but GCIQL is unbiased (Kostrikov et al., 2022; Park et al., 2023).

To extract a policy from the learned value functions, we can use either advantage-weighted regression (AWR) (Peters & Schaal, 2007; Peng et al., 2019) or behavior-constrained deep deterministic policy gradient (DDPG+BC) (Fujimoto & Gu, 2021). GCIVL uses the following value-only variant of the AWR objective (Park et al., 2023):

$$J_{\text{AWR}}^V(\pi) = \mathbb{E}_{(s,a,s') \sim p^{\mathcal{D}}(s,a,s'), g \sim p_{\text{mixed}}^{\mathcal{D}}(g|s)} [e^{\alpha(V(s', g) - V(s, g))} \log \pi(a | s, g)], \quad (5)$$

where α is the temperature hyperparameter. In our experiments, GCIQL mainly uses the following DDPG+BC objective (which is known to be better than AWR (Park et al., 2024a)):

$$J_{\text{DDPG+BC}}(\pi) = \mathbb{E}_{(s,a) \sim p^{\mathcal{D}}(s,a), g \sim p_{\text{mixed}}^{\mathcal{D}}(g|s)} [Q(s, \pi^{\mu}(s, g), g) + \alpha \log \pi(a | s, g)], \quad (6)$$

where $\pi^{\mu}(s, g) = \mathbb{E}_{a \sim \pi(a|s,g)}[a]$. In discrete-action environments, GCIQL uses the following Q version of AWR:

$$J_{\text{AWR}}^Q(\pi) = \mathbb{E}_{(s,a,s') \sim p^{\mathcal{D}}(s,a,s'), g \sim p_{\text{mixed}}^{\mathcal{D}}(g|s)} [e^{\alpha(Q(s,a,g) - V(s,g))} \log \pi(a | s, g)]. \quad (7)$$

In practice, we use standard double-value learning for GCIVL and GCIQL (Park et al., 2023), and Q normalization for DDPG+BC (Fujimoto & Gu, 2021).

Quasimetric RL (QRL). QRL (Wang et al., 2023) is a goal-conditioned value learning algorithm based on quasimetric learning, where a quasimetric means an asymmetric metric. In deterministic environments, the shortest path length between two states $d^*(s, g)$ is equivalent to the optimal undiscounted goal-conditioned value function $V^*(s, g)$ under the $(-1, 0)$ -sparse reward function (Wang et al., 2023): $V^*(s, g) = -d^*(s, g)$. The main idea of QRL is to explicitly leverage the quasimetric property (*i.e.*, triangle inequality) of shortest path lengths, namely $d^*(s, w) + d^*(w, g) \geq d^*(s, g)$ for any $s, w, g \in \mathcal{S}$, by modeling it with a quasimetric network architecture like MRN (Liu et al., 2023) or IQE (Wang & Isola, 2022). Concretely, QRL maximizes the following constrained optimization objective:

$$\text{maximize } \mathbb{E}_{s \sim p^{\mathcal{D}}(s), g \sim p_{\text{rand}}^{\mathcal{D}}(g)} [d(s, g)] \quad (8)$$

$$\text{s.t. } \mathbb{E}_{(s,s') \sim p^{\mathcal{D}}(s,s')} [(d(s, s') - 1)^2] \leq \varepsilon^2, \quad (9)$$

where $d(s, g)$ is a quasimetric distance function (*e.g.*, an IQE network) and ε is a hyperparameter that controls the strength of the constraint.

To extract a policy from the value function $V(s, g) = -d(s, g)$, QRL uses the value-only AWR loss Equation (5) in discrete-action MDPs. In continuous-action MDPs, QRL additionally fits a dynamics model $f(\phi(s), a): \mathcal{Z} \times \mathcal{A} \rightarrow \mathcal{Z}$ (Wang et al., 2023), where \mathcal{Z} denotes a latent space and $\phi(s): \mathcal{S} \rightarrow \mathcal{Z}$ denotes the representation function used in the quasimetric distance function: namely, $d(s, g) = \tilde{d}(\phi(s), \phi(g))$ (e.g., ϕ is the interval representation function in IQE). The dynamics loss is as follows:

$$\mathcal{L}_{\text{dyn}}(f) = \mathbb{E}_{(s, a, s') \sim p^{\mathcal{D}}(s, a, s')} \left[\tilde{d}(\phi(s'), f(\phi(s), a)) + \tilde{d}(f(\phi(s), a), \phi(s')) \right], \quad (10)$$

where QRL jointly trains both d and f without stop-gradients. Based on the dynamics model f , QRL maximizes the following DDPG+BC-like loss:

$$J_{\text{DDPG+BC}}(\pi) = \mathbb{E}_{(s, a) \sim p^{\mathcal{D}}(s, a), g \sim p_{\text{mixed}}^{\mathcal{D}}(g|s)} \left[-\tilde{d}(f(\phi(s), \pi^{\mu}(s, g)), \phi(g)) + \alpha \log \pi(a | s, g) \right], \quad (11)$$

where we use the same notation as Equation (6).

In practice, we employ a softplus loss shaping for the quasimetric loss and delta prediction for the dynamics model, as in Wang et al. (2023).

Contrastive RL (CRL). CRL (Eysenbach et al., 2022) is a “one-step” GCRL algorithm that first trains a Monte Carlo goal-conditioned value function using contrastive learning and performs a one-step policy improvement. CRL maximizes the following binary NCE objective (Ma & Collins, 2018) with respect to $f(s, a, g): \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$:

$$J_{\text{CRL}}(f) = \mathbb{E}_{(s, a) \sim p^{\mathcal{D}}(s, a), g \sim p_{\text{geom}}^{\mathcal{D}}(g|s, a), g^- \sim p_{\text{rand}}^{\mathcal{D}}(g)} [\log \sigma(f(s, a, g)) + \log(1 - \sigma(f(s, a, g^-)))], \quad (12)$$

where $\sigma: \mathbb{R} \rightarrow (0, 1)$ denotes the sigmoid function. The optimal solution to the above objective is given as $f(s, a, g) = \log(p_{\text{geom}}^{\mathcal{D}}(g | s, a)/p_{\text{rand}}^{\mathcal{D}}(g))$. Given the equivalence between the geometric future goal distribution ($p_{\text{geom}}^{\mathcal{D}}$) and the Monte Carlo goal-conditioned Q function (Q^{MC}) under the $(0, 1)$ -sparse reward function $r(s, g) = \mathbf{1}_{\{g\}}(s)$ (Eysenbach et al., 2022), we get the following relation: $f(s, a, g) = \log Q^{\text{MC}}(s, a, g) + C(g)$, where C is a function that only depends on g . In practice, f is parameterized as $f(s, a, g) = \phi(s, a)^{\top} \psi(g)/\sqrt{d}$ with $\phi: \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{Z} = \mathbb{R}^d$ and $\psi: \mathcal{S} \rightarrow \mathcal{Z} = \mathbb{R}^d$ (note that this inner-product parameterization is universal (Park et al., 2024c)), and we use the future goals from the other states in the same batch as g^- . We also employ the double-value learning technique for f (Eysenbach et al., 2022).

For policy extraction, in continuous-action MDPs, we employ DDPG+BC (Equation (6)) using f instead of Q . In discrete-action MDPs, we use AWR (Equation (5)) using (f, f^V) instead of (Q, V) , where we additionally train a contrastive value function $f^V(s, g): \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}$ using

$$J_{\text{CRL-V}}(f^V) = \mathbb{E}_{s \sim p^{\mathcal{D}}(s), g \sim p_{\text{geom}}^{\mathcal{D}}(g|s), g^- \sim p_{\text{rand}}^{\mathcal{D}}(g)} [\log \sigma(f^V(s, g)) + \log(1 - \sigma(f^V(s, g^-)))], \quad (13)$$

with a similar inner product parameterization for f^V .

Hierarchical implicit Q-learning (HIQL). HIQL (Park et al., 2023) is an offline GCRL algorithm that extracts two policies from a single goal-conditioned value function. HIQL first trains GCIVL (Equation (2)) with a parameterized value function defined as $V(s, g) = \tilde{V}(s, \phi(s, g))$, where $\phi: \mathcal{S} \times \mathcal{S} \rightarrow \mathcal{Z}$ serves as a (state-dependent) subgoal representation function. Based on the GCIVL value function, HIQL extracts a high-level policy $\pi^h: \mathcal{S} \times \mathcal{S} \rightarrow \Delta(\mathcal{Z})$ and a low-level policy $\pi^\ell: \mathcal{S} \times \mathcal{Z} \rightarrow \Delta(\mathcal{A})$ with the following AWR-like objectives:

$$J_{\text{HIQL}}^h(\pi^h) = \mathbb{E}_{(s_t, s_{t+k}) \sim p^{\mathcal{D}}, g \sim p_{\text{mixed}}^{\mathcal{D}}(g|s_t)} \left[e^{\alpha(V(s_{t+k}, g) - V(s_t, g))} \log \pi^h(\phi(s_t, s_{t+k}) | s_t, g) \right], \quad (14)$$

$$J_{\text{HIQL}}^\ell(\pi^\ell) = \mathbb{E}_{(s_t, a_t, s_{t+1}, s_{t+k}) \sim p^{\mathcal{D}}} \left[e^{\alpha(V(s_{t+1}, s_{t+k}) - V(s_t, s_{t+k}))} \log \pi^\ell(a_t | s_t, \phi(s_t, s_{t+k})) \right], \quad (15)$$

where we omit the arguments in $p^{\mathcal{D}}$, and k denotes a hyperparameter corresponding to the subgoal step. For simplicity, we ignore some edge cases in the objectives above (e.g., when $t + k$ exceeds the

trajectory boundary, in which case we truncate); we refer to [Park et al. \(2023\)](#) or our code for the full details. Intuitively, the high-level policy predicts the representation of the optimal k -step subgoal, and the low-level policy predicts the optimal action based on the predicted subgoal.

In practice, following [Park et al. \(2023\)](#), we use the double-value learning technique, normalize the output of ϕ , and allow gradient flows from the low-level AWR loss into ϕ (only) in pixel-based environments.

E IMPLEMENTATION DETAILS

We provide the full implementation details in this section. We include the reference implementations, environments, and data-generation scripts in the supplementary material. For datasets, during the review process, instead of releasing the full datasets (which requires a service that can host more than 100GB of files anonymously), we provide the exact commands for our data-generation scripts to reproduce all 59 datasets. After the anonymous review period ends, we will package and release the entire OGBench codebase and datasets.

E.1 TASKS AND DATASETS

In this section, we provide further information about our tasks and datasets. We provide basic specifications about the environments and datasets in Tables 6 and 7.

Locomotion tasks. For OGBench AntMaze and AntSoccer, we adopt the Ant model from D4RL AntMaze ([Fu et al., 2020](#)) (which is based on the Ant in OpenAI Gym ([Brockman et al., 2016](#); [Towers et al., 2024](#)), but with a more restricted joint range) and the soccer ball model from the DeepMind Control suite ([Tassa et al., 2018](#)). For HumanoidMaze, we adopt the Humanoid model from the DeepMind Control suite ([Tassa et al., 2018](#)).

To collect datasets, we train expert low-level directional (AntMaze and HumanoidMaze) or goal-reaching (AntSoccer) policies using SAC with dense reward functions for 400K (AntMaze), 40M (HumanoidMaze), or 12M (AntSoccer) steps. When collecting datasets, we add Gaussian noise with a standard deviation of 1.0 (explore) or 0.2 (others) to expert actions.

The success criteria for the locomotion tasks are based only on the distance between the agent (or the ball in AntSoccer) and the goal location. In particular, the tasks do not consider joint positions determining success, as in previous works ([Fu et al., 2020](#); [Park et al., 2023](#)).

Manipulation tasks. We adopt the UR5e robot arm and Robotiq 2F-85 gripper models from MuJoCo Menagerie ([Zakka et al., 2022](#)), and the drawer, window, and button box models from Meta-World ([Yu et al., 2019](#)). The robot is end-effector controlled with a 5-D action space, where the dimensions correspond to the displacements in the x position, y position, z position, gripper yaw, and gripper opening. In all manipulation tasks, we place invisible walls to prevent cubes from moving into an area beyond the robot arm’s reach. These invisible walls also prevent simply hiding the cube behind the drawer to match the goal image of the block-in-drawer task in visual-scene.

The datasets are collected by closed-loop scripted policies. When collecting datasets, we first sample the degree of action noise at the beginning of each episode, and collect a trajectory with the chosen amount of action noise. This ensures high coverage while having a sufficient number of optimal trajectories.

The success criteria for the manipulation tasks are based only on the object configurations; the arm pose is not considered when determining success. For cubes, only the distances between the goal positions and their current positions are considered, and their orientations are ignored.

Drawing tasks. We modify the original Powderworld environment ([Frans & Isola, 2023](#)) to make it offline and goal-conditioned. We provide three versions of Powderworld tasks: powderworld-easy uses two elements (plant and stone), powderworld-medium uses five elements (sand, water, fire, plant, and stone), and powderworld-hard uses eight elements (sand, water, fire, plant, stone, gas, wood, and ice). An action in Powderworld corresponds to drawing a 4×4 -sized square with a specific element brush on the 32×32 -sized board. Since naïvely implementing this atomic action requires up to 512-dimensional discrete actions, we split it into three *sequential* 8-dimensional actions that correspond to element selection, x -coordinate selection, and y -coordinate selection. To ensure

Table 5: **Dataset categories.** We list the dataset categories used to aggregate results in Figure 2.

Category	Datasets
Locomotion (states)	antmaze-medium-navigate-v0
	antmaze-large-navigate-v0
	antmaze-giant-navigate-v0
	humanoidmaze-medium-navigate-v0
	humanoidmaze-large-navigate-v0
	humanoidmaze-giant-navigate-v0
	antsoccer-arena-navigate-v0
Locomotion (pixels)	antsoccer-medium-navigate-v0
	visual-antmaze-medium-navigate-v0
	visual-antmaze-large-navigate-v0
	visual-antmaze-giant-navigate-v0
	visual-humanoidmaze-medium-navigate-v0
	visual-humanoidmaze-large-navigate-v0
	visual-humanoidmaze-giant-navigate-v0
Manipulation (states)	cube-single-play-v0
	cube-double-play-v0
	cube-triple-play-v0
	cube-quadruple-play-v0
	scene-play-v0
	puzzle-3x3-play-v0
	puzzle-4x4-play-v0
Manipulation (pixels)	puzzle-4x5-play-v0
	puzzle-4x6-play-v0
	visual-cube-single-play-v0
	visual-cube-double-play-v0
	visual-cube-triple-play-v0
	visual-cube-quadruple-play-v0
	visual-scene-play-v0
Drawing (pixels)	visual-puzzle-3x3-play-v0
	visual-puzzle-4x4-play-v0
	visual-puzzle-4x5-play-v0
Stitching (states)	visual-puzzle-4x6-play-v0
	powderworld-easy-play-v0
	powderworld-medium-play-v0
Stitching (pixels)	powderworld-hard-play-v0
	antmaze-medium-stitch-v0
	antmaze-large-stitch-v0
	antmaze-giant-stitch-v0
	humanoidmaze-medium-stitch-v0
	humanoidmaze-large-stitch-v0
	humanoidmaze-giant-stitch-v0
Exploratory (states)	antsoccer-arena-stitch-v0
	antsoccer-medium-stitch-v0
	visual-antmaze-medium-stitch-v0
Exploratory (pixels)	visual-antmaze-large-stitch-v0
	visual-antmaze-giant-stitch-v0
	visual-humanoidmaze-medium-stitch-v0
Stochastic (states)	visual-humanoidmaze-large-stitch-v0
	visual-humanoidmaze-giant-stitch-v0
Stochastic (pixels)	antmaze-medium-explore-v0
	antmaze-large-explore-v0
Stochastic (pixels)	visual-antmaze-medium-explore-v0
	visual-antmaze-large-explore-v0
Stochastic (pixels)	antmaze-teleport-navigate-v0
	visual-antmaze-teleport-navigate-v0

full observability, we add three additional dimensions that contain information about the currently selected element and x coordinate to the original $32 \times 32 \times 3$ -dimensional image, which results in a $32 \times 32 \times 6$ -dimensional observation space. When the agent selects an invalid action (which can only happen in powderworld- $\{\text{easy}, \text{medium}\}$, which has fewer than 8 elements), the environment instead uses a randomly sampled valid action.

The datasets are collected by a scripted policy that randomly draws squares and lines or fills the entire board with randomly selected brushes. With a probability of 0.5, it performs a random action (*i.e.*, places a random element on a randomly sampled position).

For the success criterion for evaluation goals, we use the following procedure to allow for some tolerance: For each pixel in the goal image, we check if the current image has a matching pixel that is

Table 6: Environment specifications.

Environment Type	Environment	State Dim.	Action Dim.	Maximum Episode Length
antmaze	antmaze-medium-v0	29	8	1000
	antmaze-large-v0	29	8	1000
	antmaze-giant-v0	29	8	1000
	antmaze-teleport-v0	29	8	1000
humanoidmaze	humanoidmaze-medium-v0	69	21	2000
	humanoidmaze-large-v0	69	21	2000
	humanoidmaze-giant-v0	69	21	4000
antsoccer	antsoccer-arena-v0	42	8	1000
	antsoccer-medium-v0	42	8	1000
visual-antmaze	visual-antmaze-medium-v0	$64 \times 64 \times 3$	8	1000
	visual-antmaze-large-v0	$64 \times 64 \times 3$	8	1000
	visual-antmaze-giant-v0	$64 \times 64 \times 3$	8	1000
	visual-antmaze-teleport-v0	$64 \times 64 \times 3$	8	1000
visual-humanoidmaze	visual-humanoidmaze-medium-v0	$64 \times 64 \times 3$	21	2000
	visual-humanoidmaze-large-v0	$64 \times 64 \times 3$	21	2000
	visual-humanoidmaze-giant-v0	$64 \times 64 \times 3$	21	4000
cube	cube-single-v0	28	5	200
	cube-double-v0	37	5	500
	cube-triple-v0	46	5	1000
	cube-quadruple-v0	55	5	1000
scene	scene-v0	40	5	750
puzzle	puzzle-3x3-v0	55	5	500
	puzzle-4x4-v0	83	5	500
	puzzle-4x5-v0	99	5	1000
	puzzle-4x6-v0	115	5	1000
visual-cube	visual-cube-single-v0	$64 \times 64 \times 3$	5	200
	visual-cube-double-v0	$64 \times 64 \times 3$	5	500
	visual-cube-triple-v0	$64 \times 64 \times 3$	5	1000
	visual-cube-quadruple-v0	$64 \times 64 \times 3$	5	1000
visual-scene	visual-scene-v0	$64 \times 64 \times 3$	5	750
visual-puzzle	visual-puzzle-3x3-v0	$64 \times 64 \times 3$	5	500
	visual-puzzle-4x4-v0	$64 \times 64 \times 3$	5	500
	visual-puzzle-4x5-v0	$64 \times 64 \times 3$	5	1000
	visual-puzzle-4x6-v0	$64 \times 64 \times 3$	5	1000
powderworld	powderworld-easy-v0	$32 \times 32 \times 6$	8 (discrete)	500
	powderworld-medium-v0	$32 \times 32 \times 6$	8 (discrete)	500
	powderworld-hard-v0	$32 \times 32 \times 6$	8 (discrete)	500

shifted by up to one pixel in any direction. We then compute the error as the number of pixels that do not match, and consider the task successful if the error is below a certain threshold.

E.2 METHODS

Our implementations of six offline GCRL algorithms (GCBC, GCIVL, GCIQL, QRL, CRL, and HIQL) are based on JAX (Bradbury et al., 2018). In our benchmark, each run typically takes 2-5 hours (state-based tasks) or 5-12 hours (pixel-based tasks) on an A5000 GPU, depending on the task and algorithm.

For benchmarking, we periodically evaluate the performance (goal success rate in percentage) of each agent on each test-time goal with 50 rollouts every 100K steps, and report the average success rate across the last three evaluation epochs (*i.e.*, at 800K, 900K, and 1M steps for state-based tasks and at 300K, 400K, and 500K steps for pixel-based tasks). That is, the performance of each agent is averaged over 750 rollouts (3 evaluation epochs \times 5 test-time goals \times 50 rollouts). While we use a relatively large number of evaluation rollouts for robustness, researchers can adjust the number of evaluation rollouts (*e.g.*, 20 episodes for each test-time goal) to reduce the computational burden.

We provide the full list of common hyperparameters in Table 8. We find that methods are more sensitive to policy extraction hyperparameters (*e.g.*, the BC coefficient in DDPG+BC) (Park et al., 2024a), and report these in a separate table (Table 9). Specifically, for each method, we use the same value learning hyperparameters across the benchmark except for the discount factor γ (Table 8), but individually tune the policy extraction hyperparameters (*e.g.*, AWR α and DDPG+BC α) for each dataset category (Tables 8 and 9).

Table 7: Dataset specifications.

Environment Type	Dataset Type	Dataset	# Transitions	# Episodes	Data Episode Length
antmaze	navigate	antmaze-medium-navigate-v0	1M	1000	1000
		antmaze-large-navigate-v0	1M	1000	1000
		antmaze-giant-navigate-v0	1M	500	2000
		antmaze-teleport-navigate-v0	1M	1000	1000
	stitch	antmaze-medium-stitch-v0	1M	5000	200
		antmaze-large-stitch-v0	1M	5000	200
		antmaze-giant-stitch-v0	1M	5000	200
		antmaze-teleport-stitch-v0	1M	5000	200
	explore	antmaze-medium-explore-v0	5M	10000	500
		antmaze-large-explore-v0	5M	10000	500
		antmaze-teleport-explore-v0	5M	10000	500
humanoidmaze	navigate	humanoidmaze-medium-navigate-v0	2M	1000	2000
		humanoidmaze-large-navigate-v0	2M	1000	2000
		humanoidmaze-giant-navigate-v0	4M	1000	4000
	stitch	humanoidmaze-medium-stitch-v0	2M	5000	400
		humanoidmaze-large-stitch-v0	2M	5000	400
		humanoidmaze-giant-stitch-v0	4M	10000	400
antsoccer	navigate	antsoccer-arena-navigate-v0	1M	1000	1000
		antsoccer-medium-navigate-v0	4M	4000	1000
	stitch	antsoccer-arena-stitch-v0	1M	5000	200
		antsoccer-medium-stitch-v0	4M	8000	500
visual-antmaze	navigate	visual-antmaze-medium-navigate-v0	1M	1000	1000
		visual-antmaze-large-navigate-v0	1M	1000	1000
		visual-antmaze-giant-navigate-v0	1M	500	2000
		visual-antmaze-teleport-navigate-v0	1M	1000	1000
	stitch	visual-antmaze-medium-stitch-v0	1M	5000	200
		visual-antmaze-large-stitch-v0	1M	5000	200
		visual-antmaze-giant-stitch-v0	1M	5000	200
		visual-antmaze-teleport-stitch-v0	1M	5000	200
	explore	visual-antmaze-medium-explore-v0	5M	10000	500
		visual-antmaze-large-explore-v0	5M	10000	500
		visual-antmaze-teleport-explore-v0	5M	10000	500
visual-humanoidmaze	navigate	visual-humanoidmaze-medium-navigate-v0	2M	1000	2000
		visual-humanoidmaze-large-navigate-v0	2M	1000	2000
		visual-humanoidmaze-giant-navigate-v0	4M	1000	4000
	stitch	visual-humanoidmaze-medium-stitch-v0	2M	5000	400
		visual-humanoidmaze-large-stitch-v0	2M	5000	400
		visual-humanoidmaze-giant-stitch-v0	4M	10000	400
cube	play	cube-single-play-v0	1M	1000	1000
		cube-double-play-v0	1M	1000	1000
		cube-triple-play-v0	3M	3000	1000
		cube-quadruple-play-v0	5M	5000	1000
scene	play	scene-play-v0	1M	1000	1000
puzzle	play	puzzle-3x3-play-v0	1M	1000	1000
		puzzle-4x4-play-v0	1M	1000	1000
		puzzle-4x5-play-v0	2M	2000	1000
		puzzle-4x6-play-v0	3M	3000	1000
visual-cube	play	visual-cube-single-play-v0	1M	1000	1000
		visual-cube-double-play-v0	1M	1000	1000
		visual-cube-triple-play-v0	3M	3000	1000
		visual-cube-quadruple-play-v0	5M	5000	1000
visual-scene	play	visual-scene-play-v0	1M	1000	1000
visual-puzzle	play	visual-puzzle-3x3-play-v0	1M	1000	1000
		visual-puzzle-4x4-play-v0	1M	1000	1000
		visual-puzzle-4x5-play-v0	2M	2000	1000
		visual-puzzle-4x6-play-v0	3M	3000	1000
powderworld	play	powderworld-easy-play-v0	1M	1000	1000
		powderworld-medium-play-v0	3M	3000	1000
		powderworld-hard-play-v0	5M	5000	1000

We apply layer normalization (Ba et al., 2016) to the value networks, but not to the policy networks. In pixel-based environments, we use a smaller version of the IMPALA encoder (Espeholt et al., 2018). We use random-crop image augmentation (with a probability of 0.5) for pixel-based manipulation tasks, but not for pixel-based locomotion or drawing tasks, as we find it to be helpful mainly on manipulation tasks. For policies, we parameterize the action distribution with a unit-variance Gaussian distribution. We find that using a fixed standard deviation is especially important for DDPG+BC.

Table 8: Common hyperparameters.

Hyperparameter	Value
Learning rate	0.0003
Optimizer	Adam (Kingma & Ba, 2015)
# gradient steps	1000000 (states), 500000 (pixels)
Minibatch size	1024 (states), 256 (pixels)
MLP dimensions	(512, 512, 512)
Nonlinearity	GELU (Hendrycks & Gimpel, 2016)
Target smoothing coefficient	0.005
Discount factor γ	0.995 (antmaze-giant, humanoidmaze), 0.99 (others)
Image augmentation probability	0.5 (pixel-based manipulation), 0 (others)
GCIVL/GCIQL expectile κ	0.9
GCIQL expectile κ	0.9
QRL quasimetric	IQE (Wang & Isola, 2022)
QRL latent dimension	512 (64 components \times 8-dimensional latents)
QRL margin ϵ	0.05
CRL latent dimension	512
HIQL expectile κ	0.7
HIQL subgoal step k	100 (humanoidmaze), 25 (other locomotion), 10 (others)
HIQL subgoal representation dimension	10
Actor ($p_{\text{cur}}^{\mathcal{D}}, p_{\text{traj}}^{\mathcal{D}}, p_{\text{geom}}^{\mathcal{D}}, p_{\text{rand}}^{\mathcal{D}}$) ratio for $p_{\text{mixed}}^{\mathcal{D}}$	(0, 0.5, 0, 0.5) (stitch), (0, 0, 0, 1) (explore), (0, 1, 0, 0) (others)
Value ($p_{\text{cur}}^{\mathcal{D}}, p_{\text{traj}}^{\mathcal{D}}, p_{\text{geom}}^{\mathcal{D}}, p_{\text{rand}}^{\mathcal{D}}$) ratio for $p_{\text{mixed}}^{\mathcal{D}}$	(0.2, 0, 0.5, 0.3)

Table 9: Hyperparameters for policy extraction. Each cell indicates the policy extraction method and its α value (*i.e.*, the temperature (AWR) or the BC coefficient (DDPG+BC)).

Environment Type	Dataset Type	Dataset	GCIVL	GCIQL	QRL	CRL	HIQL
antmaze	navigate	antmaze-medium-navigate-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
		antmaze-large-navigate-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
		antmaze-giant-navigate-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
		antmaze-teleport-navigate-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
	stitch	antmaze-medium-stitch-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
		antmaze-large-stitch-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
		antmaze-giant-stitch-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
		antmaze-teleport-stitch-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
	explore	antmaze-medium-explore-v0	AWR 10.0	DDPG 0.01	DDPG 0.001	DDPG 0.003	AWR 10.0
		antmaze-large-explore-v0	AWR 10.0	DDPG 0.01	DDPG 0.001	DDPG 0.003	AWR 10.0
		antmaze-teleport-explore-v0	AWR 10.0	DDPG 0.01	DDPG 0.001	DDPG 0.003	AWR 10.0
humanoidmaze	navigate	humanoidmaze-medium-navigate-v0	AWR 10.0	DDPG 0.1	DDPG 0.001	DDPG 0.1	AWR 3.0
		humanoidmaze-large-navigate-v0	AWR 10.0	DDPG 0.1	DDPG 0.001	DDPG 0.1	AWR 3.0
		humanoidmaze-giant-navigate-v0	AWR 10.0	DDPG 0.1	DDPG 0.001	DDPG 0.1	AWR 3.0
	stitch	humanoidmaze-medium-stitch-v0	AWR 10.0	DDPG 0.1	DDPG 0.001	DDPG 0.1	AWR 3.0
		humanoidmaze-large-stitch-v0	AWR 10.0	DDPG 0.1	DDPG 0.001	DDPG 0.1	AWR 3.0
		humanoidmaze-giant-stitch-v0	AWR 10.0	DDPG 0.1	DDPG 0.001	DDPG 0.1	AWR 3.0
antsoccer	navigate	antsoccer-arena-navigate-v0	AWR 10.0	DDPG 0.1	DDPG 0.003	DDPG 0.3	AWR 3.0
		antsoccer-medium-navigate-v0	AWR 10.0	DDPG 0.1	DDPG 0.003	DDPG 0.3	AWR 3.0
	stitch	antsoccer-arena-stitch-v0	AWR 10.0	DDPG 0.1	DDPG 0.003	DDPG 0.3	AWR 3.0
		antsoccer-medium-stitch-v0	AWR 10.0	DDPG 0.1	DDPG 0.003	DDPG 0.3	AWR 3.0
visual-antmaze	navigate	visual-antmaze-medium-navigate-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
		visual-antmaze-large-navigate-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
		visual-antmaze-giant-navigate-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
		visual-antmaze-teleport-navigate-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
	stitch	visual-antmaze-medium-stitch-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
		visual-antmaze-large-stitch-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
		visual-antmaze-giant-stitch-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
		visual-antmaze-teleport-stitch-v0	AWR 10.0	DDPG 0.3	DDPG 0.003	DDPG 0.1	AWR 3.0
	explore	visual-antmaze-medium-explore-v0	AWR 10.0	DDPG 0.01	DDPG 0.001	DDPG 0.003	AWR 10.0
		visual-antmaze-large-explore-v0	AWR 10.0	DDPG 0.01	DDPG 0.001	DDPG 0.003	AWR 10.0
		visual-antmaze-teleport-explore-v0	AWR 10.0	DDPG 0.01	DDPG 0.001	DDPG 0.003	AWR 10.0
visual-humanoidmaze	navigate	visual-humanoidmaze-medium-navigate-v0	AWR 10.0	DDPG 0.1	DDPG 0.001	DDPG 0.1	AWR 3.0
		visual-humanoidmaze-large-navigate-v0	AWR 10.0	DDPG 0.1	DDPG 0.001	DDPG 0.1	AWR 3.0
		visual-humanoidmaze-giant-navigate-v0	AWR 10.0	DDPG 0.1	DDPG 0.001	DDPG 0.1	AWR 3.0
	stitch	visual-humanoidmaze-medium-stitch-v0	AWR 10.0	DDPG 0.1	DDPG 0.001	DDPG 0.1	AWR 3.0
		visual-humanoidmaze-large-stitch-v0	AWR 10.0	DDPG 0.1	DDPG 0.001	DDPG 0.1	AWR 3.0
		visual-humanoidmaze-giant-stitch-v0	AWR 10.0	DDPG 0.1	DDPG 0.001	DDPG 0.1	AWR 3.0
cube	play	cube-single-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
		cube-double-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
		cube-triple-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
		cube-quadruple-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
scene	play	scene-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
puzzle	play	puzzle-3x3-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
		puzzle-4x4-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
		puzzle-4x5-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
		puzzle-4x6-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
visual-cube	play	visual-cube-single-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
		visual-cube-double-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
		visual-cube-triple-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
		visual-cube-quadruple-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
visual-scene	play	visual-scene-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
visual-puzzle	play	visual-puzzle-3x3-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
		visual-puzzle-4x4-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
		visual-puzzle-4x5-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
		visual-puzzle-4x6-play-v0	AWR 10.0	DDPG 0.03	DDPG 0.03	DDPG 0.1	AWR 3.0
powderworld	play	powderworld-easy-play-v0	AWR 3.0	AWR 3.0	AWR 3.0	AWR 3.0	AWR 3.0
		powderworld-medium-play-v0	AWR 3.0	AWR 3.0	AWR 3.0	AWR 3.0	AWR 3.0
		powderworld-hard-play-v0	AWR 3.0	AWR 3.0	AWR 3.0	AWR 3.0	AWR 3.0

F EVALUATION GOALS AND PER-GOAL BENCHMARKING RESULTS

Each task in OGBench provides five evaluation goals. We provide their full image descriptions in Figures 4 to 10, and the full per-goal evaluation results in Tables 10 to 16, which share the same layout as Table 2.




Figure 4: AntMaze and HumanoidMaze goals.




Figure 5: AntSoccer goals.

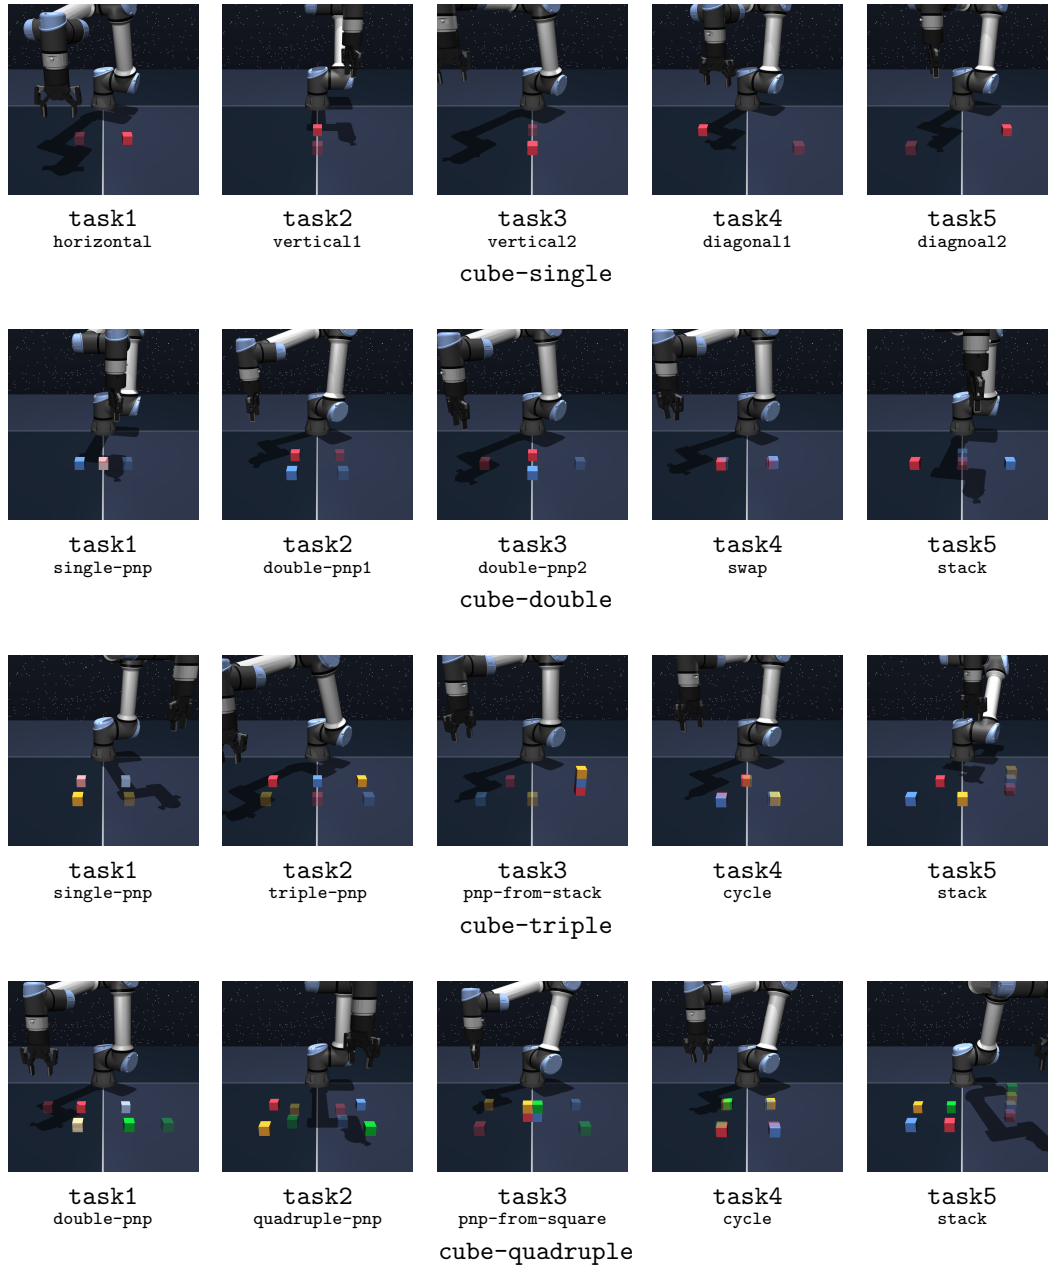


Figure 6: Cube goals.

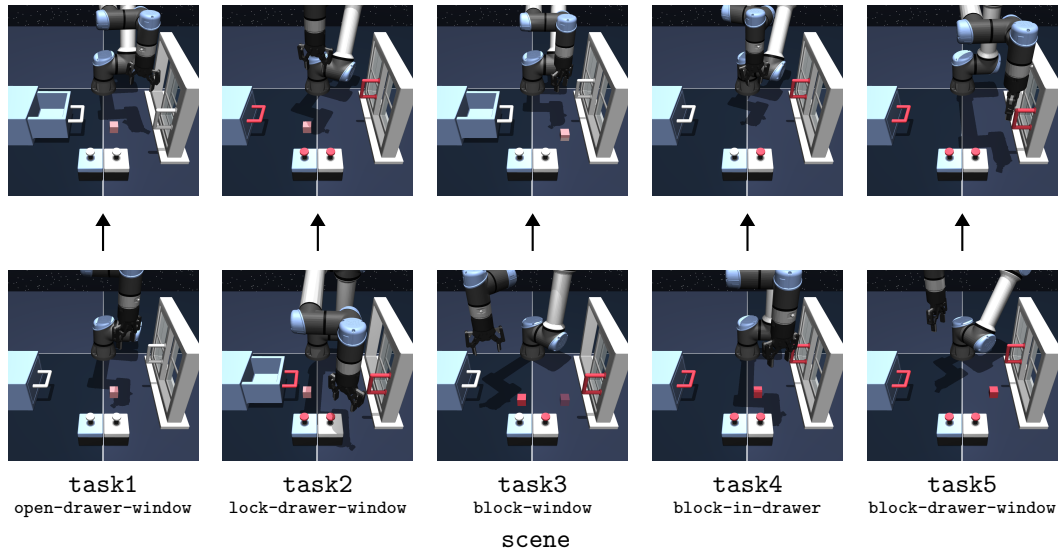


Figure 7: Scene goals.

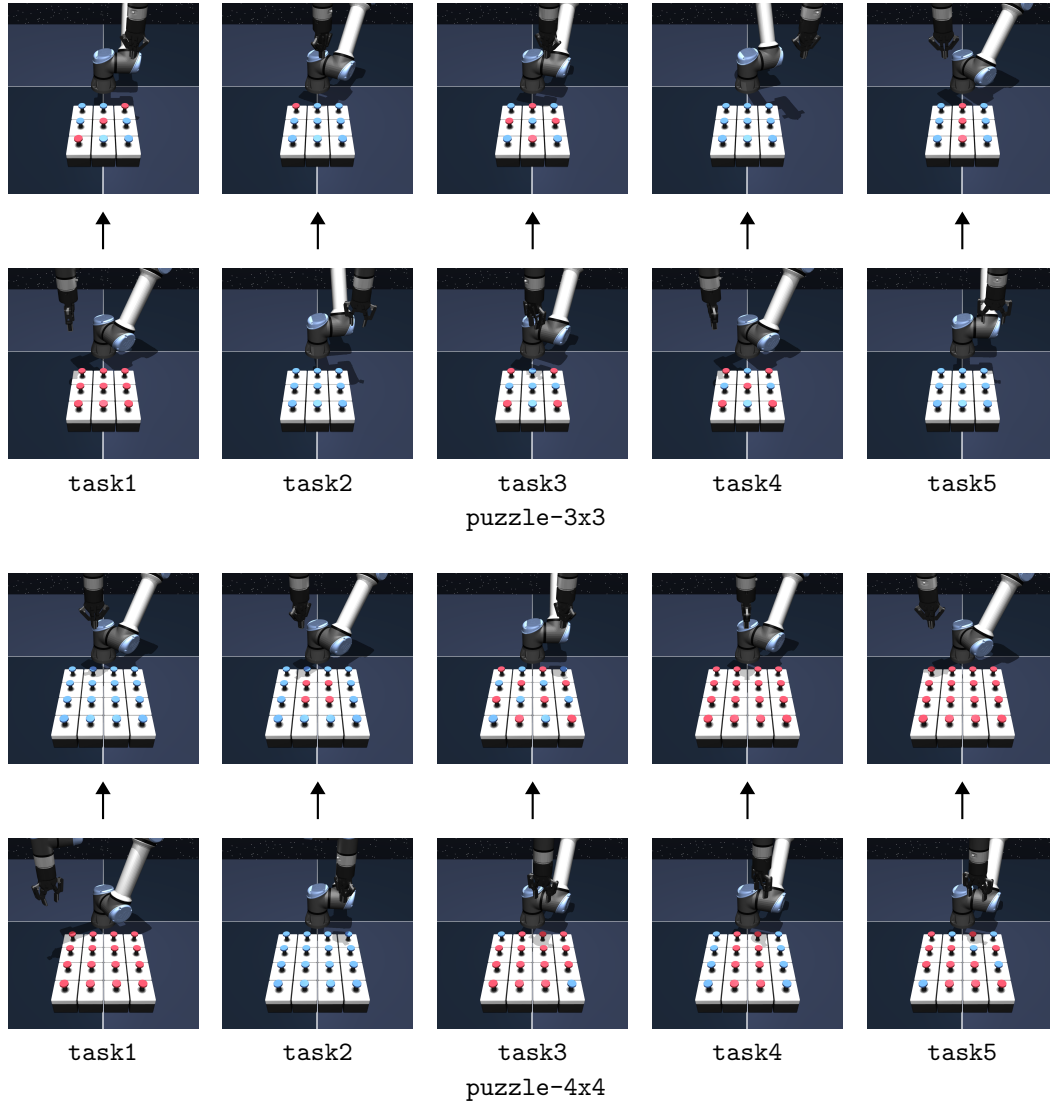
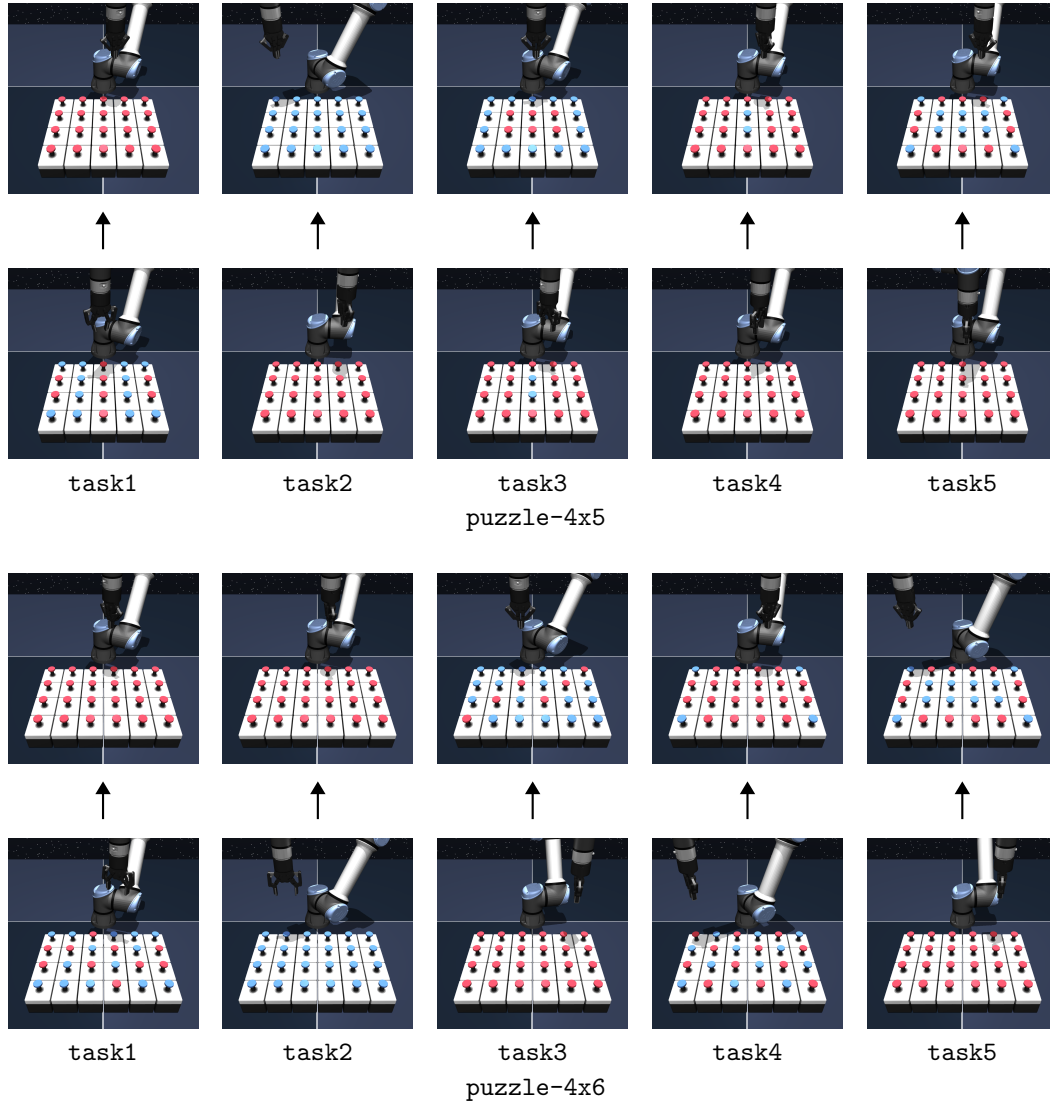


Figure 8: **Puzzle goals.**

Figure 9: **Puzzle goals.**

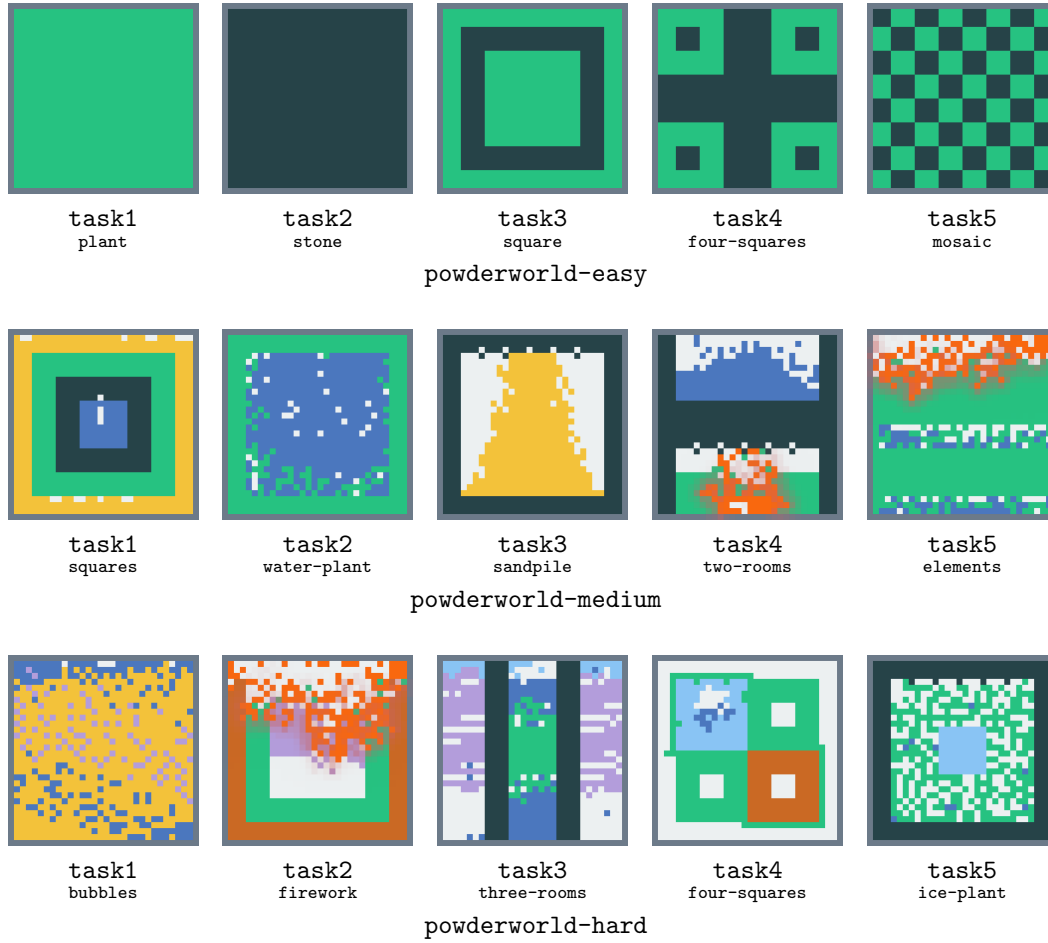


Figure 10: Powderworld goals.

Table 10: Full results on AntMaze.

Environment Type	Dataset Type	Dataset	Task	GCBC	GCIVL	GCIQL	QRL	CRL	HIQL
antmaze	navigate	antmaze-medium-navigate-v0	task1	35.2 \pm 9.4	80.8 \pm 10.4	63.3 \pm 9.2	93.0 \pm 2.2	97.0 \pm 1.3	94.4 \pm 2.2
			task2	21.0 \pm 6.9	85.0 \pm 4.8	77.5 \pm 7.8	89.8 \pm 5.4	94.7 \pm 2.0	96.8 \pm 1.2
			task3	24.3 \pm 5.7	60.5 \pm 12.6	70.9 \pm 8.3	85.6 \pm 6.4	92.4 \pm 2.8	96.2 \pm 2.4
			task4	27.9 \pm 7.0	42.0 \pm 25.4	59.2 \pm 11.9	83.3 \pm 3.9	94.2 \pm 4.5	96.2 \pm 2.4
			task5	36.7 \pm 9.8	91.8 \pm 2.7	85.0 \pm 7.0	87.6 \pm 7.6	96.0 \pm 1.9	96.1 \pm 1.5
			overall	29.0 \pm 4.0	72.0 \pm 7.9	71.2 \pm 4.1	87.9 \pm 3.5	94.8 \pm 1.3	95.9 \pm 0.9
		antmaze-large-navigate-v0	task1	6.0 \pm 3.3	16.5 \pm 12.2	20.8 \pm 5.9	70.8 \pm 15.3	90.6 \pm 2.8	92.8 \pm 2.9
			task2	16.4 \pm 3.5	5.2 \pm 5.5	25.2 \pm 6.7	77.2 \pm 6.5	61.9 \pm 13.8	78.1 \pm 8.7
			task3	65.2 \pm 4.2	49.2 \pm 18.0	79.6 \pm 5.3	94.2 \pm 2.3	90.9 \pm 2.5	96.2 \pm 1.7
			task4	13.9 \pm 2.6	2.5 \pm 2.4	18.7 \pm 5.8	63.9 \pm 8.5	85.1 \pm 11.4	94.2 \pm 2.1
			task5	18.5 \pm 4.1	5.0 \pm 2.5	26.5 \pm 9.2	66.8 \pm 9.2	85.2 \pm 3.0	93.8 \pm 3.0
			overall	24.0 \pm 1.8	15.7 \pm 5.4	34.2 \pm 3.6	74.6 \pm 6.4	82.8 \pm 4.0	91.0 \pm 2.2
		antmaze-giant-navigate-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	1.5 \pm 2.0	2.1 \pm 2.3	47.0 \pm 10.0
			task2	0.1 \pm 0.2	0.0 \pm 0.0	0.2 \pm 0.3	17.3 \pm 5.0	21.2 \pm 9.9	74.2 \pm 4.0
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	13.7 \pm 7.5	4.8 \pm 4.6	54.7 \pm 7.2
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	18.2 \pm 6.4	34.7 \pm 9.1	68.8 \pm 4.9
			task5	0.8 \pm 0.6	1.2 \pm 1.2	1.0 \pm 0.8	17.7 \pm 4.7	15.6 \pm 10.3	81.5 \pm 4.0
			overall	0.2 \pm 0.2	0.2 \pm 0.2	0.2 \pm 0.1	13.7 \pm 3.4	15.7 \pm 3.3	65.2 \pm 5.3
		antmaze-teleport-navigate-v0	task1	16.9 \pm 4.9	34.8 \pm 5.2	26.1 \pm 4.7	31.2 \pm 5.7	35.2 \pm 5.3	37.2 \pm 5.4
			task2	50.6 \pm 4.9	40.9 \pm 5.3	57.6 \pm 7.9	46.9 \pm 22.3	91.8 \pm 3.1	66.2 \pm 8.4
			task3	21.8 \pm 3.3	36.0 \pm 8.2	31.2 \pm 4.8	34.8 \pm 6.1	46.6 \pm 4.4	37.1 \pm 4.5
			task4	24.7 \pm 4.5	44.7 \pm 2.8	33.3 \pm 5.5	33.1 \pm 5.9	49.7 \pm 1.9	29.9 \pm 2.0
			task5	13.6 \pm 6.0	38.0 \pm 5.9	25.8 \pm 8.6	28.5 \pm 8.2	43.7 \pm 2.8	41.2 \pm 8.3
			overall	25.5 \pm 2.8	38.9 \pm 3.0	34.8 \pm 5.4	34.9 \pm 5.3	53.4 \pm 1.6	42.3 \pm 2.5
	stitch	antmaze-medium-stitch-v0	task1	70.5 \pm 32.8	76.2 \pm 12.5	17.1 \pm 12.0	42.7 \pm 20.4	43.3 \pm 9.9	92.2 \pm 2.0
			task2	65.2 \pm 19.2	80.5 \pm 4.5	22.5 \pm 16.3	60.7 \pm 12.1	46.3 \pm 13.5	93.9 \pm 3.1
			task3	21.3 \pm 15.4	16.4 \pm 12.4	41.2 \pm 9.1	71.9 \pm 28.9	45.8 \pm 17.4	95.2 \pm 2.0
			task4	0.8 \pm 2.1	0.0 \pm 0.0	31.9 \pm 8.9	79.6 \pm 9.5	52.7 \pm 19.3	93.1 \pm 2.0
			task5	69.6 \pm 32.9	47.1 \pm 19.9	33.6 \pm 13.7	40.8 \pm 18.3	75.1 \pm 8.4	94.6 \pm 2.7
			overall	45.5 \pm 11.0	44.0 \pm 5.6	29.3 \pm 6.2	59.1 \pm 6.7	52.6 \pm 6.2	93.8 \pm 5.9
		antmaze-large-stitch-v0	task1	1.9 \pm 2.2	23.4 \pm 9.1	0.2 \pm 0.3	6.9 \pm 4.6	0.8 \pm 0.9	85.1 \pm 4.9
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	9.8 \pm 4.9	4.4 \pm 3.8	23.9 \pm 16.2
			task3	15.2 \pm 14.4	69.0 \pm 5.6	37.1 \pm 10.1	73.3 \pm 7.7	43.1 \pm 11.4	94.1 \pm 3.2
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.2	0.6 \pm 1.2	4.8 \pm 4.9	70.2 \pm 7.9
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	1.2 \pm 1.0	0.7 \pm 1.9	59.9 \pm 9.5
			overall	3.4 \pm 2.8	18.5 \pm 2.4	7.5 \pm 2.0	18.4 \pm 2.1	10.8 \pm 1.8	66.7 \pm 5.3
		antmaze-giant-stitch-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.3 \pm 0.5
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.2	0.0 \pm 0.0	5.3 \pm 4.8
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.5
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	2.6 \pm 3.1
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.2	2.0 \pm 2.4	0.0 \pm 0.0	0.3 \pm 0.7
			overall	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.4 \pm 0.5	0.0 \pm 0.0	1.8 \pm 1.6
		antmaze-teleport-stitch-v0	task1	20.6 \pm 13.5	39.2 \pm 7.4	12.5 \pm 4.2	21.8 \pm 5.9	30.1 \pm 5.8	44.2 \pm 5.1
			task2	38.7 \pm 11.5	43.7 \pm 6.0	17.9 \pm 6.9	21.8 \pm 6.3	30.0 \pm 4.3	42.4 \pm 3.0
			task3	34.4 \pm 11.7	35.6 \pm 7.9	18.4 \pm 3.9	24.8 \pm 7.2	23.2 \pm 10.8	26.4 \pm 3.8
			task4	46.5 \pm 6.2	44.0 \pm 3.6	17.8 \pm 4.6	23.8 \pm 9.3	37.6 \pm 3.7	26.1 \pm 3.5
			task5	16.4 \pm 14.0	33.1 \pm 6.3	17.0 \pm 5.6	26.3 \pm 4.5	31.9 \pm 7.5	40.3 \pm 5.5
			overall	31.3 \pm 6.3	39.1 \pm 3.0	16.7 \pm 2.3	23.7 \pm 5.2	30.5 \pm 3.6	35.9 \pm 2.5
	explore	antmaze-medium-explore-v0	task1	2.8 \pm 5.7	10.1 \pm 7.5	11.8 \pm 5.6	0.7 \pm 0.8	1.5 \pm 1.5	29.2 \pm 16.8
			task2	1.2 \pm 2.3	73.8 \pm 8.9	52.9 \pm 8.4	0.6 \pm 1.2	7.7 \pm 6.4	83.8 \pm 10.2
			task3	1.0 \pm 1.6	0.0 \pm 0.0	0.2 \pm 0.3	3.3 \pm 4.7	4.2 \pm 6.5	18.5 \pm 23.7
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.5	0.0 \pm 0.0	0.1 \pm 0.2
			task5	3.2 \pm 3.6	10.0 \pm 5.9	0.1 \pm 0.2	0.8 \pm 0.9	2.3 \pm 1.7	52.2 \pm 27.4
			overall	1.6 \pm 1.0	18.8 \pm 2.7	13.0 \pm 2.1	1.1 \pm 1.3	3.1 \pm 2.1	36.7 \pm 10.5
		antmaze-large-explore-v0	task1	0.0 \pm 0.0	37.1 \pm 12.2	0.9 \pm 1.0	0.0 \pm 0.0	0.2 \pm 0.7	1.2 \pm 2.8
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task3	0.2 \pm 0.5	12.0 \pm 5.9	1.1 \pm 1.4	0.1 \pm 0.2	1.3 \pm 1.1	18.2 \pm 24.5
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			overall	0.0 \pm 0.1	9.8 \pm 3.4	0.4 \pm 0.3	0.0 \pm 0.0	0.3 \pm 0.2	3.9 \pm 5.0
		antmaze-teleport-explore-v0	task1	2.2 \pm 2.3	31.5 \pm 4.2	0.5 \pm 0.6	0.2 \pm 0.3	2.1 \pm 1.5	32.2 \pm 11.5
			task2	0.2 \pm 0.3	2.3 \pm 4.2	7.6 \pm 5.7	0.3 \pm 0.9	4.9 \pm 3.9	32.7 \pm 17.3
			task3	3.5 \pm 3.3	47.8 \pm 2.9	13.1 \pm 7.6	3.6 \pm 3.8	46.8 \pm 5.9	34.2 \pm 16.4
			task4	2.2 \pm 2.0	47.2 \pm 4.8	13.5 \pm 8.4	4.5 \pm 4.0	15.9 \pm 11.9	36.9 \pm 18.6
			task5	3.8 \pm 1.9	31.1 \pm 1.9	1.7 \pm 1.4	2.9 \pm 3.2	27.8 \pm 5.0	34.2 \pm 14.3
			overall	2.4 \pm 1.1	32.0 \pm 1.7	7.3 \pm 3.3	2.3 \pm 2.1	19.5 \pm 2.2	34.1 \pm 15.0

Table 11: Full results on HumanoidMaze and AntSoccer.

Environment Type	Dataset Type	Dataset	Task	GCBC	GCIVL	GCIQL	QRL	CRL	HIQL
humanoidmaze	navigate	humanoidmaze-medium-navigate-v0	task1	4.5 \pm 1.2	22.5 \pm 4.7	22.6 \pm 6.1	11.9 \pm 7.1	83.9 \pm 3.0	95.1 \pm 2.0
			task2	8.3 \pm 3.8	42.2 \pm 7.6	49.1 \pm 6.0	25.3 \pm 8.4	80.0 \pm 5.2	96.3 \pm 2.1
			task3	11.8 \pm 3.3	14.8 \pm 3.0	12.4 \pm 6.3	25.2 \pm 10.5	43.4 \pm 10.6	79.2 \pm 5.9
			task4	1.7 \pm 0.9	0.1 \pm 0.2	0.9 \pm 0.5	15.6 \pm 6.6	4.8 \pm 5.1	75.2 \pm 6.1
			task5	11.5 \pm 3.9	40.3 \pm 7.8	51.3 \pm 7.6	28.9 \pm 12.5	87.3 \pm 6.6	96.8 \pm 1.5
			overall	7.5 \pm 1.6	24.0 \pm 2.2	27.3 \pm 2.5	21.4 \pm 8.2	59.9 \pm 3.8	88.5 \pm 1.8
		humanoidmaze-large-navigate-v0	task1	0.8 \pm 0.8	5.8 \pm 1.8	3.2 \pm 1.9	2.8 \pm 1.6	36.2 \pm 10.8	66.8 \pm 4.3
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.4 \pm 0.5	0.2 \pm 0.3	2.2 \pm 2.7
			task3	2.6 \pm 0.9	5.5 \pm 1.7	4.6 \pm 2.0	17.0 \pm 5.7	54.3 \pm 17.0	88.1 \pm 2.9
			task4	2.1 \pm 1.2	0.2 \pm 0.3	1.1 \pm 0.9	4.5 \pm 1.9	22.9 \pm 10.5	42.2 \pm 10.5
			task5	1.3 \pm 0.9	0.8 \pm 0.8	0.9 \pm 0.8	1.8 \pm 1.4	6.5 \pm 3.6	47.3 \pm 10.5
			overall	1.4 \pm 0.4	2.5 \pm 0.6	2.0 \pm 0.7	5.3 \pm 1.3	24.0 \pm 4.2	49.3 \pm 4.2
		humanoidmaze-giant-navigate-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.9 \pm 0.9	13.0 \pm 7.3
			task2	0.4 \pm 0.5	0.8 \pm 1.2	1.4 \pm 1.0	1.8 \pm 0.7	8.6 \pm 4.7	34.8 \pm 10.6
			task3	0.1 \pm 0.2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	2.0 \pm 1.7	10.9 \pm 3.6
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.2	3.2 \pm 2.3	2.0 \pm 2.0
			task5	0.6 \pm 0.6	0.2 \pm 0.5	0.9 \pm 1.0	1.7 \pm 0.9	1.0 \pm 1.4	1.8 \pm 1.6
			overall	0.2 \pm 0.1	0.2 \pm 0.3	0.5 \pm 0.3	0.7 \pm 0.3	3.1 \pm 1.5	12.5 \pm 4.2
	stitch	humanoidmaze-medium-stitch-v0	task1	20.3 \pm 7.0	13.2 \pm 3.4	12.3 \pm 3.2	6.4 \pm 4.8	26.7 \pm 6.7	83.9 \pm 5.0
			task2	48.8 \pm 12.4	6.7 \pm 1.8	7.9 \pm 5.1	12.8 \pm 5.5	36.8 \pm 6.5	93.8 \pm 2.1
			task3	23.8 \pm 8.5	25.2 \pm 3.3	20.0 \pm 6.6	30.3 \pm 5.6	40.5 \pm 3.9	85.8 \pm 4.0
			task4	2.8 \pm 2.5	1.0 \pm 1.0	2.5 \pm 1.6	18.3 \pm 5.2	27.8 \pm 6.9	85.7 \pm 3.7
			task5	49.4 \pm 8.2	15.6 \pm 3.0	17.7 \pm 7.1	22.1 \pm 2.3	49.2 \pm 4.6	90.2 \pm 3.7
			overall	29.0 \pm 4.8	12.3 \pm 1.7	12.1 \pm 3.1	18.0 \pm 2.1	36.2 \pm 2.5	87.9 \pm 2.2
		humanoidmaze-large-stitch-v0	task1	3.4 \pm 4.1	1.7 \pm 0.7	0.7 \pm 0.6	0.2 \pm 0.3	0.4 \pm 0.5	20.7 \pm 5.1
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.3	0.1 \pm 0.2	5.3 \pm 2.3
			task3	20.4 \pm 11.5	2.6 \pm 2.1	1.0 \pm 1.1	16.0 \pm 7.0	12.7 \pm 2.5	84.2 \pm 3.9
			task4	1.5 \pm 0.8	0.6 \pm 0.7	0.5 \pm 0.6	0.9 \pm 0.7	3.9 \pm 1.1	19.3 \pm 3.5
			task5	2.5 \pm 1.6	1.0 \pm 1.1	0.2 \pm 0.5	0.0 \pm 0.0	2.7 \pm 1.3	12.2 \pm 2.3
			overall	5.6 \pm 2.9	1.2 \pm 0.5	0.5 \pm 0.2	3.5 \pm 1.4	3.9 \pm 0.6	28.4 \pm 2.6
		humanoidmaze-giant-stitch-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	1.3 \pm 1.7
			task2	0.2 \pm 0.3	0.7 \pm 0.5	0.2 \pm 0.5	0.6 \pm 0.7	0.2 \pm 0.3	11.8 \pm 6.1
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	2.2 \pm 1.7
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.2	0.7 \pm 0.9
			task5	0.1 \pm 0.2	0.2 \pm 0.5	0.6 \pm 0.6	1.2 \pm 0.9	0.4 \pm 0.7	0.4 \pm 0.8
			overall	0.1 \pm 0.1	0.2 \pm 0.1	0.2 \pm 0.1	0.4 \pm 0.3	0.2 \pm 0.1	3.3 \pm 1.9
antsoccer	navigate	antsoccer-arena-navigate-v0	task1	7.2 \pm 2.8	60.8 \pm 3.5	60.0 \pm 5.7	12.4 \pm 2.1	32.0 \pm 5.3	67.2 \pm 3.9
			task2	9.7 \pm 2.9	44.6 \pm 5.5	55.8 \pm 4.2	8.2 \pm 3.6	26.7 \pm 4.0	58.8 \pm 4.3
			task3	3.1 \pm 1.6	62.2 \pm 5.5	63.2 \pm 5.6	12.4 \pm 2.5	28.4 \pm 5.2	76.2 \pm 3.7
			task4	3.0 \pm 0.9	23.0 \pm 5.5	27.9 \pm 6.3	2.8 \pm 1.7	10.8 \pm 2.2	30.0 \pm 3.5
			task5	2.9 \pm 1.1	42.5 \pm 5.5	42.4 \pm 7.1	2.5 \pm 3.1	14.8 \pm 3.7	55.8 \pm 4.4
			overall	5.2 \pm 0.6	46.6 \pm 2.6	49.9 \pm 2.4	7.6 \pm 1.8	22.6 \pm 2.2	57.6 \pm 2.4
		antsoccer-medium-navigate-v0	task1	9.4 \pm 2.6	17.0 \pm 4.6	28.8 \pm 5.4	7.5 \pm 7.7	13.0 \pm 4.3	44.6 \pm 6.2
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.2	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.3
			task3	0.2 \pm 0.3	0.0 \pm 0.0	0.5 \pm 0.9	0.7 \pm 0.8	0.5 \pm 0.6	2.2 \pm 1.2
			task4	0.4 \pm 0.5	0.1 \pm 0.2	0.7 \pm 0.8	0.7 \pm 0.8	0.8 \pm 0.7	3.0 \pm 0.8
			task5	0.7 \pm 0.6	2.6 \pm 1.3	3.9 \pm 3.0	0.7 \pm 1.9	2.0 \pm 1.2	12.8 \pm 5.2
			overall	2.1 \pm 0.5	3.9 \pm 1.0	6.8 \pm 1.2	1.9 \pm 2.0	3.2 \pm 1.0	12.5 \pm 1.5
	stitch	antsoccer-arena-stitch-v0	task1	72.9 \pm 5.3	36.6 \pm 3.8	5.7 \pm 1.6	1.6 \pm 1.2	2.3 \pm 1.6	23.8 \pm 2.4
			task2	36.1 \pm 19.1	12.8 \pm 4.3	2.3 \pm 1.2	1.6 \pm 1.8	0.8 \pm 0.5	14.1 \pm 4.4
			task3	6.4 \pm 15.0	34.2 \pm 8.8	1.0 \pm 0.9	0.6 \pm 0.6	0.3 \pm 0.4	19.9 \pm 2.9
			task4	6.8 \pm 12.4	11.2 \pm 3.3	0.2 \pm 0.5	0.0 \pm 0.0	0.1 \pm 0.2	6.7 \pm 2.0
			task5	0.1 \pm 0.2	12.1 \pm 3.2	1.3 \pm 0.9	0.1 \pm 0.2	0.1 \pm 0.2	12.1 \pm 5.3
			overall	24.5 \pm 7.8	21.4 \pm 3.1	2.1 \pm 0.4	0.8 \pm 0.5	0.7 \pm 0.3	15.3 \pm 1.3
		antsoccer-medium-stitch-v0	task1	10.0 \pm 6.9	4.2 \pm 1.9	0.2 \pm 0.5	0.1 \pm 0.2	0.2 \pm 0.3	21.2 \pm 6.0
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.2
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.2
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.3 \pm 0.9
			overall	2.0 \pm 1.4	0.8 \pm 0.4	0.0 \pm 0.1	0.0 \pm 0.0	0.1 \pm 0.1	4.3 \pm 1.2

Table 12: Full results on Visual AntMaze.

Environment Type	Dataset Type	Dataset	Task	GCBC	GCIVL	GCIQL	QRL	CRL	HIQL
visual-antmaze	navigate	visual-antmaze-medium-navigate-v0	task1	17.2 \pm 5.8	30.2 \pm 7.2	15.8 \pm 3.0	0.0 \pm 0.0	91.5 \pm 2.2	89.8 \pm 3.9
			task2	7.5 \pm 2.3	20.7 \pm 6.3	7.0 \pm 1.6	0.0 \pm 0.0	94.3 \pm 1.8	91.7 \pm 7.0
			task3	17.0 \pm 1.4	23.5 \pm 4.7	15.8 \pm 3.6	0.0 \pm 0.0	97.8 \pm 1.5	93.5 \pm 3.5
			task4	11.8 \pm 2.1	21.2 \pm 3.3	8.7 \pm 2.2	0.0 \pm 0.0	94.3 \pm 1.7	94.2 \pm 1.6
			task5	3.7 \pm 2.5	16.3 \pm 4.9	6.5 \pm 1.9	0.0 \pm 0.0	93.5 \pm 2.2	94.2 \pm 4.7
			overall	11.4 \pm 1.5	22.4 \pm 2.4	10.8 \pm 1.4	0.0 \pm 0.0	94.3 \pm 0.6	92.7 \pm 3.8
		visual-antmaze-large-navigate-v0	task1	3.2 \pm 1.5	6.8 \pm 2.3	4.2 \pm 2.6	0.2 \pm 0.3	77.5 \pm 4.6	59.8 \pm 10.2
			task2	3.5 \pm 2.7	4.2 \pm 0.6	2.2 \pm 1.3	0.0 \pm 0.0	80.2 \pm 3.0	28.2 \pm 8.5
			task3	4.3 \pm 1.8	6.2 \pm 1.6	4.2 \pm 1.5	0.7 \pm 1.3	90.5 \pm 2.5	84.8 \pm 9.5
			task4	4.2 \pm 1.8	5.3 \pm 3.0	5.5 \pm 1.1	0.5 \pm 0.6	88.5 \pm 2.6	45.8 \pm 7.1
			task5	4.5 \pm 1.7	4.8 \pm 0.8	3.7 \pm 2.2	0.0 \pm 0.0	82.7 \pm 2.1	44.3 \pm 9.7
			overall	3.9 \pm 0.5	5.5 \pm 1.0	3.9 \pm 0.8	0.3 \pm 0.4	83.9 \pm 1.2	52.6 \pm 8.6
		visual-antmaze-giant-navigate-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	16.7 \pm 2.2	1.7 \pm 1.3
			task2	0.7 \pm 0.5	2.0 \pm 1.4	0.7 \pm 0.5	0.2 \pm 0.3	73.0 \pm 8.8	11.7 \pm 7.7
			task3	0.0 \pm 0.0	0.2 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	22.0 \pm 6.1	2.2 \pm 2.8
			task4	0.3 \pm 0.7	0.3 \pm 0.7	0.0 \pm 0.0	0.0 \pm 0.0	47.2 \pm 5.1	3.7 \pm 2.5
			task5	1.2 \pm 0.8	2.3 \pm 2.6	0.7 \pm 0.5	0.3 \pm 0.7	77.3 \pm 4.7	13.3 \pm 11.0
			overall	0.4 \pm 0.3	1.0 \pm 0.8	0.3 \pm 0.2	0.1 \pm 0.2	47.2 \pm 1.7	6.5 \pm 4.3
		visual-antmaze-teleport-navigate-v0	task1	1.8 \pm 1.8	5.5 \pm 0.8	1.8 \pm 1.3	2.8 \pm 2.5	31.8 \pm 3.3	31.7 \pm 4.7
			task2	5.5 \pm 2.7	9.2 \pm 3.3	8.8 \pm 1.8	5.8 \pm 3.9	73.3 \pm 7.8	40.5 \pm 6.1
			task3	8.8 \pm 1.0	12.5 \pm 2.6	9.0 \pm 2.5	9.7 \pm 4.1	46.7 \pm 3.2	33.3 \pm 0.9
			task4	10.0 \pm 1.7	10.5 \pm 2.3	8.2 \pm 3.3	6.0 \pm 3.7	50.5 \pm 4.2	44.5 \pm 5.0
			task5	1.2 \pm 0.6	3.3 \pm 1.4	3.0 \pm 1.3	4.0 \pm 2.2	36.0 \pm 4.8	32.7 \pm 6.6
			overall	5.5 \pm 0.8	8.2 \pm 1.2	6.2 \pm 1.2	5.7 \pm 2.7	47.7 \pm 2.2	36.5 \pm 1.5
	stitch	visual-antmaze-medium-stitch-v0	task1	79.5 \pm 3.7	0.3 \pm 0.7	0.2 \pm 0.3	0.0 \pm 0.0	33.2 \pm 4.4	74.8 \pm 8.3
			task2	89.8 \pm 4.0	1.2 \pm 1.6	0.2 \pm 0.3	0.0 \pm 0.0	69.0 \pm 5.3	85.3 \pm 7.4
			task3	69.0 \pm 17.7	15.0 \pm 5.8	8.3 \pm 1.2	0.0 \pm 0.0	88.0 \pm 1.1	92.0 \pm 1.4
			task4	0.8 \pm 0.8	7.2 \pm 3.9	2.7 \pm 0.8	0.3 \pm 0.7	70.2 \pm 11.8	88.2 \pm 3.7
			task5	97.2 \pm 1.1	5.7 \pm 3.2	0.7 \pm 0.5	0.2 \pm 0.3	84.8 \pm 5.0	93.3 \pm 1.4
			overall	67.3 \pm 3.6	5.9 \pm 2.1	2.4 \pm 0.3	0.1 \pm 0.2	69.0 \pm 2.2	86.7 \pm 2.4
		visual-antmaze-large-stitch-v0	task1	26.0 \pm 10.7	0.2 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	5.5 \pm 1.5	35.5 \pm 4.6
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	1.5 \pm 0.8	3.3 \pm 2.0
			task3	73.3 \pm 14.0	2.8 \pm 2.4	0.2 \pm 0.3	2.0 \pm 1.6	35.7 \pm 10.2	87.2 \pm 6.3
			task4	7.3 \pm 5.1	0.8 \pm 0.6	0.2 \pm 0.3	0.7 \pm 0.9	7.8 \pm 1.1	7.3 \pm 4.5
			task5	11.2 \pm 4.8	0.0 \pm 0.0	0.0 \pm 0.0	0.3 \pm 0.4	5.0 \pm 2.2	6.0 \pm 1.2
			overall	23.6 \pm 2.7	0.8 \pm 0.6	0.1 \pm 0.1	0.6 \pm 0.6	11.1 \pm 2.5	27.9 \pm 2.1
		visual-antmaze-giant-stitch-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task2	1.3 \pm 2.2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.7 \pm 0.5
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.5 \pm 0.6	0.2 \pm 0.3
			overall	0.3 \pm 0.4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.1	0.2 \pm 0.2
		visual-antmaze-teleport-stitch-v0	task1	37.2 \pm 4.3	1.8 \pm 1.8	1.3 \pm 1.2	0.5 \pm 0.3	20.3 \pm 5.2	36.0 \pm 5.5
			task2	36.5 \pm 2.7	1.8 \pm 1.5	0.8 \pm 0.6	1.0 \pm 1.2	39.5 \pm 8.8	38.0 \pm 3.0
			task3	17.3 \pm 5.6	2.0 \pm 0.5	2.0 \pm 1.2	2.8 \pm 3.9	31.8 \pm 8.8	35.7 \pm 5.0
			task4	39.2 \pm 8.7	0.8 \pm 0.8	0.2 \pm 0.3	1.8 \pm 2.8	44.7 \pm 7.4	37.2 \pm 5.7
			task5	28.7 \pm 1.4	0.7 \pm 1.3	0.7 \pm 0.9	1.0 \pm 1.2	22.2 \pm 9.2	38.3 \pm 5.1
			overall	31.8 \pm 3.1	1.4 \pm 0.8	1.0 \pm 0.4	1.4 \pm 1.6	31.7 \pm 6.4	37.0 \pm 3.5
	explore	visual-antmaze-medium-explore-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.3	0.0 \pm 0.0	0.3 \pm 0.4
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.8 \pm 1.7
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0
			overall	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.1	0.0 \pm 0.0	0.2 \pm 0.4
		visual-antmaze-large-explore-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			overall	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
		visual-antmaze-teleport-explore-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task3	0.0 \pm 0.0	0.3 \pm 0.7	0.3 \pm 0.7	0.0 \pm 0.0	3.5 \pm 0.5	38.0 \pm 8.2
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	28.5 \pm 12.9
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.5 \pm 0.3	0.3 \pm 0.4	1.8 \pm 1.8	27.2 \pm 17.5
			overall	0.0 \pm 0.0	0.1 \pm 0.1	0.2 \pm 0.2	0.1 \pm 0.1	1.0 \pm 0.4	18.7 \pm 7.6

Table 13: Full results on Visual HumanoidMaze.

Environment Type	Dataset Type	Dataset	Task	GCBC	GCIVL	GCIQL	QRL	CRL	HIQL
visual-humanoidmaze	navigate	visual-humanoidmaze-medium-navigate-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.3	0.0 \pm 0.0
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.3	0.0 \pm 0.0
			task3	0.0 \pm 0.0	0.2 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	1.7 \pm 0.7	0.5 \pm 0.6
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	2.7 \pm 1.8	0.3 \pm 0.7
			overall	0.0 \pm 0.0	0.0 \pm 0.1	0.0 \pm 0.0	0.0 \pm 0.0	0.9 \pm 0.4	0.2 \pm 0.3
		visual-humanoidmaze-large-navigate-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			overall	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
		visual-humanoidmaze-giant-navigate-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			overall	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
	stitch	visual-humanoidmaze-medium-stitch-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task2	0.2 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task3	3.0 \pm 0.9	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	3.3 \pm 1.8	1.0 \pm 1.6
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.3	0.0 \pm 0.0
			task5	1.2 \pm 0.6	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.3 \pm 0.4	0.0 \pm 0.0
			overall	0.9 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.8 \pm 0.4	0.2 \pm 0.3
		visual-humanoidmaze-large-stitch-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			overall	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
		visual-humanoidmaze-giant-stitch-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			overall	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0

Table 14: Full results on Cube, Scene, and Puzzle.

Environment Type	Dataset Type	Dataset	Task	GCBC	GCIVL	GCIQL	QRL	CRL	HQL
cube	play	cube-single-play-v0	task1	6.2 \pm 3.1	70.0 \pm 9.5	99.6 \pm 0.5	17.1 \pm 8.4	40.4 \pm 6.0	44.7 \pm 10.3
			task2	5.2 \pm 3.0	74.8 \pm 8.5	99.8 \pm 0.3	23.2 \pm 8.9	41.6 \pm 5.0	39.5 \pm 7.4
			task3	1.7 \pm 1.6	65.2 \pm 11.0	99.9 \pm 0.2	5.2 \pm 3.6	46.4 \pm 7.3	41.9 \pm 10.8
			task4	13.6 \pm 4.3	69.8 \pm 6.1	98.6 \pm 1.4	41.2 \pm 15.1	35.9 \pm 5.8	42.5 \pm 5.8
			task5	16.8 \pm 6.0	71.8 \pm 13.9	99.3 \pm 0.8	36.5 \pm 15.4	37.8 \pm 5.6	41.4 \pm 9.9
			overall	8.7 \pm 1.6	70.3 \pm 5.3	99.4 \pm 0.4	24.7 \pm 8.1	40.4 \pm 4.6	42.0 \pm 7.5
		cube-double-play-v0	task1	8.7 \pm 5.6	57.8 \pm 10.7	54.2 \pm 10.0	21.0 \pm 7.7	12.6 \pm 1.9	9.3 \pm 3.5
			task2	0.0 \pm 0.0	8.1 \pm 3.6	16.9 \pm 8.4	0.0 \pm 0.0	0.2 \pm 0.3	0.6 \pm 1.0
			task3	0.2 \pm 0.3	1.2 \pm 1.3	8.4 \pm 4.0	0.0 \pm 0.0	0.2 \pm 0.7	0.6 \pm 0.8
			task4	0.1 \pm 0.2	3.4 \pm 1.3	11.8 \pm 4.8	0.1 \pm 0.2	0.1 \pm 0.2	0.3 \pm 0.4
			task5	0.1 \pm 0.2	6.2 \pm 2.3	17.6 \pm 5.6	0.2 \pm 0.3	0.2 \pm 0.3	0.2 \pm 0.3
			overall	1.8 \pm 1.1	15.3 \pm 2.7	21.8 \pm 5.0	4.3 \pm 1.5	2.7 \pm 0.5	2.2 \pm 0.9
		cube-triple-play-v0	task1	12.2 \pm 6.1	45.2 \pm 12.1	10.2 \pm 3.8	10.5 \pm 7.6	9.8 \pm 3.8	8.2 \pm 3.7
			task2	0.0 \pm 0.0	0.2 \pm 0.3	0.9 \pm 1.6	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task3	0.0 \pm 0.0	1.5 \pm 1.0	0.1 \pm 0.2	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.5
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			overall	2.4 \pm 1.2	9.4 \pm 2.6	2.2 \pm 0.9	2.1 \pm 1.5	2.0 \pm 0.8	1.7 \pm 0.7
		cube-quadruple-play-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.5
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			overall	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.1
scene	play	scene-play-v0	task1	7.2 \pm 3.8	73.5 \pm 10.6	42.6 \pm 6.3	10.3 \pm 5.4	4.2 \pm 3.0	74.0 \pm 4.4
			task2	0.0 \pm 0.0	42.0 \pm 7.0	60.5 \pm 14.3	0.2 \pm 0.3	0.1 \pm 0.2	33.5 \pm 8.9
			task3	0.9 \pm 0.6	51.2 \pm 17.3	61.3 \pm 12.6	0.5 \pm 0.6	0.5 \pm 0.5	33.2 \pm 13.5
			task4	0.0 \pm 0.0	14.2 \pm 7.1	0.2 \pm 0.3	15.2 \pm 6.3	0.2 \pm 0.3	29.8 \pm 3.5
			task5	0.0 \pm 0.0	10.1 \pm 2.7	22.0 \pm 8.7	0.0 \pm 0.0	0.0 \pm 0.0	8.1 \pm 2.7
			overall	1.6 \pm 0.7	38.2 \pm 4.4	37.3 \pm 6.8	5.2 \pm 1.4	1.0 \pm 0.6	35.7 \pm 4.6
		puzzle-3x3-play-v0	task1	4.1 \pm 2.2	91.6 \pm 6.9	100.0 \pm 0.0	0.2 \pm 0.3	83.1 \pm 7.4	73.7 \pm 9.4
			task2	0.1 \pm 0.2	47.8 \pm 33.0	93.8 \pm 4.6	0.0 \pm 0.0	30.8 \pm 10.7	56.7 \pm 12.9
			task3	0.0 \pm 0.0	30.7 \pm 21.0	98.8 \pm 1.4	0.0 \pm 0.0	24.7 \pm 8.2	52.7 \pm 12.6
			task4	0.0 \pm 0.0	26.8 \pm 21.9	96.2 \pm 2.0	0.0 \pm 0.0	18.8 \pm 7.0	49.5 \pm 10.6
			task5	0.1 \pm 0.2	38.5 \pm 23.8	89.8 \pm 4.6	0.0 \pm 0.0	21.7 \pm 6.1	54.8 \pm 11.8
			overall	0.9 \pm 0.5	47.1 \pm 20.8	95.7 \pm 1.1	0.0 \pm 0.1	35.8 \pm 6.3	57.4 \pm 11.1
		puzzle-4x4-play-v0	task1	0.0 \pm 0.0	46.9 \pm 4.4	45.9 \pm 6.4	0.0 \pm 0.0	0.0 \pm 0.0	18.0 \pm 4.6
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	18.6 \pm 6.8
			task3	0.0 \pm 0.0	32.1 \pm 8.0	65.0 \pm 20.1	0.0 \pm 0.0	0.0 \pm 0.0	15.5 \pm 3.6
			task4	0.0 \pm 0.0	12.7 \pm 6.3	23.2 \pm 13.6	0.0 \pm 0.0	0.0 \pm 0.0	13.1 \pm 4.9
			task5	0.0 \pm 0.0	10.6 \pm 6.2	15.5 \pm 7.4	0.0 \pm 0.0	0.0 \pm 0.0	12.9 \pm 2.4
			overall	0.0 \pm 0.0	20.4 \pm 4.0	29.9 \pm 7.3	0.0 \pm 0.0	0.0 \pm 0.0	15.6 \pm 3.8
puzzle	play	puzzle-4x5-play-v0	task1	0.0 \pm 0.0	3.1 \pm 2.4	1.3 \pm 1.6	0.0 \pm 0.0	0.0 \pm 0.0	4.0 \pm 1.8
			task2	0.0 \pm 0.0	0.3 \pm 0.7	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.3 \pm 0.5
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.7 \pm 0.8
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.2
			overall	0.0 \pm 0.0	0.7 \pm 0.6	0.3 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	1.0 \pm 0.4
		puzzle-4x6-play-v0	task1	0.0 \pm 0.0	77.8 \pm 5.7	81.8 \pm 14.2	0.0 \pm 0.0	0.0 \pm 0.0	5.3 \pm 3.7
			task2	0.0 \pm 0.0	6.1 \pm 4.4	4.2 \pm 3.3	0.0 \pm 0.0	0.0 \pm 0.0	3.2 \pm 2.0
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.1 \pm 0.2
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			overall	0.0 \pm 0.0	16.8 \pm 1.5	17.2 \pm 3.3	0.0 \pm 0.0	0.0 \pm 0.0	1.7 \pm 0.9

Table 15: Full results on Visual Cube, Visual Scene, and Visual Puzzle.

Environment Type	Dataset Type	Dataset	Task	GCBC	GCIVL	GCIQL	QRL	CRL	HIQL
visual-cube	play	visual-cube-single-play-v0	task1	10.5 \pm 3.8	88.0 \pm 4.7	79.8 \pm 6.4	12.2 \pm 19.1	65.2 \pm 26.0	99.0 \pm 1.6
			task2	27.2 \pm 7.9	70.0 \pm 15.8	0.0 \pm 0.0	3.5 \pm 3.6	58.7 \pm 5.0	99.3 \pm 0.5
			task3	5.5 \pm 1.5	72.2 \pm 7.0	90.2 \pm 2.8	33.5 \pm 29.2	53.7 \pm 24.9	94.5 \pm 0.6
			task4	19.0 \pm 4.5	73.2 \pm 9.1	27.3 \pm 9.4	2.8 \pm 1.1	59.7 \pm 18.7	99.7 \pm 0.4
			task5	24.2 \pm 10.6	71.2 \pm 4.4	13.2 \pm 7.2	4.7 \pm 0.8	60.2 \pm 12.1	99.5 \pm 0.6
			overall	17.3 \pm 4.2	74.9 \pm 1.7	42.1 \pm 2.9	11.3 \pm 9.9	59.5 \pm 16.1	98.4 \pm 0.5
		visual-cube-double-play-v0	task1	12.3 \pm 5.5	60.3 \pm 4.9	37.2 \pm 5.4	34.0 \pm 11.2	15.0 \pm 5.3	97.2 \pm 1.5
			task2	3.8 \pm 2.9	5.5 \pm 2.2	2.5 \pm 1.0	11.3 \pm 15.4	3.0 \pm 1.4	88.3 \pm 5.1
			task3	1.2 \pm 0.8	1.2 \pm 1.4	3.2 \pm 1.3	0.0 \pm 0.0	0.3 \pm 0.7	55.5 \pm 6.1
			task4	0.2 \pm 0.3	0.8 \pm 0.6	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	14.8 \pm 9.1
			task5	0.8 \pm 0.8	0.0 \pm 0.0	0.3 \pm 0.4	2.3 \pm 2.1	0.5 \pm 0.3	35.8 \pm 8.5
			overall	3.7 \pm 1.2	13.6 \pm 1.5	8.6 \pm 1.0	9.5 \pm 4.4	3.8 \pm 1.4	58.3 \pm 2.7
		visual-cube-triple-play-v0	task1	86.0 \pm 2.6	80.8 \pm 5.2	53.8 \pm 3.3	53.8 \pm 16.9	84.8 \pm 6.4	98.3 \pm 0.4
			task2	1.3 \pm 0.8	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.8 \pm 0.6	9.2 \pm 4.5
			task3	0.3 \pm 0.7	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.8 \pm 1.3	17.0 \pm 9.0
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			overall	17.5 \pm 0.8	16.2 \pm 1.0	10.8 \pm 0.7	10.8 \pm 3.4	17.3 \pm 1.6	24.9 \pm 1.0
		visual-cube-quadruple-play-v0	task1	39.8 \pm 5.0	1.2 \pm 1.4	1.3 \pm 1.2	9.2 \pm 12.6	35.8 \pm 7.2	79.8 \pm 9.8
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	2.5 \pm 2.0
			task3	0.2 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.3	2.5 \pm 1.0
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.3
			overall	8.0 \pm 1.0	0.2 \pm 0.3	0.3 \pm 0.2	1.8 \pm 2.5	7.2 \pm 1.4	17.0 \pm 1.6
visual-scene	play	visual-scene-play-v0	task1	44.8 \pm 7.5	71.7 \pm 8.9	23.5 \pm 5.3	8.0 \pm 4.2	46.8 \pm 15.6	93.8 \pm 3.5
			task2	0.0 \pm 0.0	6.8 \pm 4.9	9.8 \pm 8.7	0.0 \pm 0.0	0.8 \pm 1.0	56.2 \pm 14.8
			task3	0.3 \pm 0.4	34.0 \pm 11.9	12.2 \pm 5.4	0.5 \pm 1.0	3.5 \pm 1.6	81.7 \pm 4.8
			task4	1.5 \pm 1.4	0.2 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	1.0 \pm 1.2	30.3 \pm 11.3
			task5	0.0 \pm 0.0	0.3 \pm 0.7	0.3 \pm 0.4	0.0 \pm 0.0	0.0 \pm 0.0	41.7 \pm 7.0
			overall	9.3 \pm 1.7	22.6 \pm 1.7	9.2 \pm 2.0	1.7 \pm 1.0	10.4 \pm 3.8	60.7 \pm 5.6
		visual-puzzle-3x3-play-v0	task1	0.7 \pm 0.5	100.0 \pm 0.0	100.0 \pm 0.0	0.0 \pm 0.0	3.2 \pm 5.5	100.0 \pm 0.0
			task2	0.0 \pm 0.0	0.0 \pm 0.0	0.7 \pm 0.9	0.0 \pm 0.0	0.0 \pm 0.0	61.2 \pm 22.2
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.8 \pm 0.8	0.0 \pm 0.0	0.0 \pm 0.0	61.5 \pm 17.4
			task4	0.0 \pm 0.0	0.0 \pm 0.0	1.7 \pm 1.3	0.0 \pm 0.0	0.0 \pm 0.0	53.8 \pm 16.3
			task5	0.0 \pm 0.0	0.3 \pm 0.7	11.0 \pm 5.6	0.0 \pm 0.0	0.0 \pm 0.0	68.7 \pm 12.5
			overall	0.1 \pm 0.1	20.1 \pm 0.1	22.8 \pm 1.3	0.0 \pm 0.0	0.6 \pm 1.1	69.0 \pm 12.4
visual-puzzle	play	visual-puzzle-4x4-play-v0	task1	5.5 \pm 3.9	93.2 \pm 0.8	81.2 \pm 5.0	0.0 \pm 0.0	8.5 \pm 1.5	98.5 \pm 1.1
			task2	10.0 \pm 1.2	0.3 \pm 0.7	0.5 \pm 0.6	0.0 \pm 0.0	12.0 \pm 0.9	52.7 \pm 17.3
			task3	3.0 \pm 1.3	91.0 \pm 4.8	76.0 \pm 3.1	0.0 \pm 0.0	5.3 \pm 2.9	96.3 \pm 3.3
			task4	3.2 \pm 1.9	35.0 \pm 7.8	46.5 \pm 16.1	0.0 \pm 0.0	2.8 \pm 2.1	94.8 \pm 3.0
			task5	5.2 \pm 2.5	15.8 \pm 7.7	35.2 \pm 11.2	0.0 \pm 0.0	5.3 \pm 6.0	95.2 \pm 1.0
			overall	5.4 \pm 1.9	47.1 \pm 2.1	47.9 \pm 4.5	0.0 \pm 0.0	6.8 \pm 1.7	87.5 \pm 5.0
		visual-puzzle-4x5-play-v0	task1	1.8 \pm 1.4	0.0 \pm 0.0	1.3 \pm 1.2	0.0 \pm 0.0	1.2 \pm 0.6	6.8 \pm 2.4
			task2	0.3 \pm 0.7	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.8 \pm 0.3	0.0 \pm 0.0
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			overall	0.4 \pm 0.4	0.0 \pm 0.0	0.3 \pm 0.2	0.0 \pm 0.0	0.4 \pm 0.2	1.4 \pm 0.5
		visual-puzzle-4x6-play-v0	task1	13.7 \pm 8.2	82.3 \pm 7.3	85.3 \pm 4.4	0.0 \pm 0.0	7.7 \pm 5.2	53.0 \pm 16.2
			task2	0.8 \pm 0.3	0.7 \pm 0.8	2.2 \pm 1.0	0.0 \pm 0.0	0.0 \pm 0.0	10.8 \pm 10.3
			task3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			overall	2.9 \pm 1.7	16.6 \pm 1.6	17.5 \pm 0.9	0.0 \pm 0.0	1.5 \pm 1.0	12.8 \pm 5.1

Table 16: Full results on Powderworld.

Environment Type	Dataset Type	Dataset	Task	GCBC	GCIVL	GCIQL	QRL	CRL	HIQL
powderworld	play	powderworld-easy-play-v0	task1	0.8 \pm 1.0	80.8 \pm 13.8	94.0 \pm 4.5	36.3 \pm 15.0	40.0 \pm 27.2	45.5 \pm 31.5
			task2	0.2 \pm 0.3	73.2 \pm 17.2	97.7 \pm 1.6	17.3 \pm 6.1	31.2 \pm 21.2	39.2 \pm 30.7
			task3	0.0 \pm 0.0	98.0 \pm 2.4	99.8 \pm 0.3	10.2 \pm 5.5	11.5 \pm 13.3	33.5 \pm 28.6
			task4	0.0 \pm 0.0	98.8 \pm 1.4	99.8 \pm 0.3	7.7 \pm 4.4	7.2 \pm 8.4	4.8 \pm 3.9
			task5	0.0 \pm 0.0	97.8 \pm 1.9	97.2 \pm 3.5	3.8 \pm 3.1	2.2 \pm 3.5	0.3 \pm 0.4
			overall	0.2 \pm 0.3	89.7 \pm 4.7	97.7 \pm 1.3	15.1 \pm 6.5	18.4 \pm 12.8	24.7 \pm 18.0
		powderworld-medium-play-v0	task1	1.3 \pm 1.4	74.5 \pm 8.5	1.2 \pm 1.1	2.5 \pm 2.2	0.2 \pm 0.3	32.2 \pm 10.6
			task2	2.3 \pm 1.4	54.2 \pm 31.2	1.3 \pm 2.2	6.7 \pm 4.5	1.0 \pm 1.3	19.3 \pm 9.6
			task3	4.3 \pm 5.1	93.3 \pm 13.3	71.5 \pm 16.2	2.7 \pm 2.0	5.3 \pm 3.0	73.5 \pm 8.6
			task4	0.0 \pm 0.0	60.5 \pm 12.3	0.5 \pm 0.6	0.0 \pm 0.0	0.0 \pm 0.0	3.3 \pm 2.8
			task5	0.0 \pm 0.0	0.7 \pm 0.9	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.3	0.5 \pm 1.0
			overall	1.6 \pm 0.9	56.6 \pm 6.6	14.9 \pm 2.9	2.4 \pm 1.2	1.3 \pm 0.9	25.8 \pm 2.4
		powderworld-hard-play-v0	task1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task2	0.0 \pm 0.0	3.3 \pm 2.4	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	3.5 \pm 4.6
			task3	0.0 \pm 0.0	6.7 \pm 7.5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	2.2 \pm 3.9
			task4	0.0 \pm 0.0	7.7 \pm 8.5	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
			task5	0.5 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.3
			overall	0.1 \pm 0.1	3.5 \pm 2.3	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	1.2 \pm 1.7