

# Physically Embedded Planning Problems: New Challenges for Reinforcement Learning

Mehdi Mirza<sup>1,\*</sup>, Andrew Jaegle<sup>1,\*</sup>, Jonathan J. Hunt<sup>3,\*</sup>, Arthur Guez<sup>1,\*</sup>, Saran Tunyasuvunakool<sup>1</sup>, Alistair Muldal<sup>1</sup>, Théophane Weber<sup>1</sup>, Peter Karkus<sup>2,3</sup>, Sébastien Racanière<sup>1</sup>, Lars Buesing<sup>1</sup>, Timothy Lillicrap<sup>1</sup> and Nicolas Heess<sup>1</sup>

<sup>\*</sup>Equal contributions, <sup>1</sup>DeepMind, <sup>2</sup>National University of Singapore, <sup>3</sup>Work done at DeepMind




Figure 1 | Three symbolic planning domains (top row) are embedded in a physical environment, simulated in MuJoCo, to obtain our three proposed challenge domains (bottom row). An agent must coordinate perception, reasoning, and motor control over long time horizons to be successful on these tasks.

Recent work in deep reinforcement learning (RL) has produced algorithms capable of mastering challenging games such as Go, chess, or shogi. In these works the RL agent directly observes the natural state of the game and controls that state directly with its actions. However, when humans play such games, they do not just reason about the moves but also interact with their physical environment. They understand the state of the game by looking at the physical board in front of them and modify it by manipulating pieces using touch and fine-grained motor control. Mastering complicated physical systems with abstract goals is a central challenge for artificial intelligence, but it remains out of reach for existing RL algorithms. To encourage progress towards this goal we introduce a set of physically embedded planning problems and make them [publicly available](#). We embed challenging symbolic tasks (Sokoban, tic-tac-toe, and Go) in a physics engine to produce a set of tasks that require perception, reasoning, and motor control over long time horizons. Although existing RL algorithms can tackle the symbolic versions of these tasks, we find that they struggle to master even the simplest of their physically embedded counterparts. As a first step towards characterizing the space of solution to these tasks, we introduce a strong baseline that uses a pre-trained expert game player to provide hints in the abstract space to an RL agent’s policy while training it on the full sensorimotor control task. The resulting agent solves many of the tasks, underlining the need for methods that bridge the gap between abstract planning and embodied control. See illustrating [video](#).

## 1. Introduction

Early work in artificial intelligence (AI) often focused on symbolic reasoning and ignored the challenges of physical control and grounding, which were sometimes discounted as less central to the problem of intelligence (Brooks, 1990). Since then, many approaches to physical control problems have been developed in different disciplines, including optimal control, robotics, and reinforcement learning, striking different trade-offs between learning and human engineering.

In recent years deep reinforcement learning (RL) has shown promising results on a number of problems that exemplify different aspects of artificial and embodied intelligence. For instance, results on difficult board games such as chess, Go (Silver et al., 2018), and Sokoban (Racanière et al., 2017) have demonstrated that RL can solve tasks which, due to their combinatorial structure and the irreversibility of decisions, are often considered as *reasoning* or *planning* domains. In contrast, results such as those of (Heess et al., 2017; Hwangbo et al., 2019; Levine et al., 2016; Merel et al., 2020; OpenAI et al., 2018) demonstrate that RL algorithms can be successfully applied to problems that involve sophisticated *low-level control* of physically embodied systems, whether in manipulation or locomotion tasks that require the short-term coordination of high-dimensional bodies. RL methods have also been able to tackle domains where *perception* is a challenge, due to partial observability or a complex sensory stream (Beattie et al., 2016; Gregor et al., 2019; Paine et al., 2020; Song et al., 2020; Wayne et al., 2018).

However, many real world problems, or ethologically relevant tasks, require a combination of high-level reasoning, low-level motor control, and perception at very different temporal scales. This is well recognized in the robotics literature where long-horizon sequential reasoning has long been studied (e.g. Kambhampati et al., 1991). Through the use of hand-crafted and structural methods (which often rely on the use of privileged information), this research program has led to many impressive results, including on challenging multi-object manipulation tasks (e.g., Dantam et al., 2018; Lagriffoul et al., 2018) and on games such as chess (Dantam et al., 2011). In contrast, learning-based methods have so far paid relatively little attention to problems that combine all of the challenges mentioned above. Although progress has been made along each individual axes of complexity, domain agnostic RL methods are not competitive on long-horizon embodied planning problems against the well-engineered, domain-specific approaches common in the robotics literature.

The aim of this work is to facilitate the study of such problems with learning based approaches. We provide a set of simulation domains that combine the difficulties of planning with those of low-level motor control. This is achieved by embedding existing planning tasks (tic-tac-toe, Go, Sokoban) into a physical setting where, in order to execute a move in the game, the agent is required to control a physical body over many simulation time-steps. For example, placing a single tic-tac-toe piece requires reaching the board with a 9 DoF (degree of freedom) arm and touching the appropriate position on the board. Although learning to play tic-tac-toe and executing goal-directed reaching movements with a robotic arm are individually well within the capabilities of current deep RL techniques, combining these two control problems into a single task can pose challenges. This effect is stronger for more difficult planning tasks (e.g. going from tic-tac-toe to Go) and for problems that require more sophisticated perception or low-level control (e.g., egocentric vision or controlling a physical body with many DoFs).

Our tasks complement popular motor control tasks that are currently used as benchmarks by the RL community (e.g. Brockman et al., 2016; Tassa et al., 2018). These tasks often focus on the short-term motor control component in the form of locomotion or manipulation, but require only limited high-level reasoning. They tend to be reactive with a small effective decision horizon, and the solution can often be memorized once the right behavior is found. In contrast, our domains force the agent to respond to a combinatorial number of situations, by achieving a sequence of well-defined abstract subgoals, which need to be chosen while taking into account their long-term, and possibly irreversible, consequences. The combination of motor control and high-level reasoning demands an agent that can effectively learn to act at multiple time-scales. Our benchmark is designed with the goal of providing a fertile experimental setting for the development of RL methods for physically-grounded sequential reasoning and control, rather than to develop methods that can be directly ported to physical robots. In this respect, our benchmark is complementary to others that strike a different balance between physical and abstract reasoning complexity by more realistically simulating laboratory robotic platforms (Dantam et al., 2018; Lagriffoul et al., 2018).

Our tasks embed abstract planning problems into a physical environment with the need for motor control. This distinguishes them from tasks in which the physical environment and the reasoning problem are more intimately coupled (e.g. [Fazeli et al., 2019](#); [Lee et al., 2019b](#)). Although this design choice limits the scope of the tasks we propose, it also has several key advantages: 1) the properties and requirements of the underlying reasoning task are relatively well understood, 2) the tasks possess natural abstractions which we can exploit, 3) we know RL methods that can tackle both the high-level and low-level problems in isolation, 4) we can vary the difficulty of the underlying reasoning task while keeping the control difficulty roughly constant, and conversely 5) we can vary the difficulty of the sensorimotor RL task (and how tightly it corresponds to the abstract task) while keeping the difficulty of the underlying reasoning task constant.

As a reference, and to demonstrate the general feasibility of our tasks, we include a preliminary investigation of training RL agents on our domains in Section 4. We study vanilla model-free agents as well as agents that have access to the output of a pre-specified abstract planner (similar to the approach used by [Akkaya et al., 2019](#)). This study demonstrates that these domains remain challenging even when considerable privileged information about the structure of the problem is available. The insights gained from these experiments may inform the search for methods that can solve such domains efficiently.

## 2. Environments

We introduce a set of three discrete planning problems embedded in continuous physical domains, illustrated in Figure 1. These environments are open-sourced [here](#) and code snippets are available in Appendix Section A.




Figure 2 | **(Left pair)** MuJoBan without grid pegs, example of the top-down view and the corresponding agent’s egocentric view next to it. **(Right pair)**: Same camera views for MuJoBan with grid pegs.

### 2.1. MuJoBan

MuJoBan is a single player puzzle-solving game embedded in the MuJoCo simulation environment ([Todorov et al., 2012](#)). The puzzle is based on the 2D game of Sokoban (see [here](#)), where an agent situated on a grid has to push boxes onto target locations. Only one box can be pushed at a time, and boxes can only be pushed but not pulled. This rule leads to many actions being irreversible (e.g. pushing a box into a corner), so an agent must choose its actions carefully and deliberately to solve the game. This is why Sokoban has proven popular as an RL benchmark for studying planning-type methods ([Guez et al., 2019](#); [Junghanns and Schaeffer, 1997](#); [Quack et al., 2015](#); [Racanière et al., 2017](#)).

In MuJoBan, an embodied agent is situated in the 3D maze equivalent of a Sokoban level. The visual appearance of the environment (such as color of walls) is varied from episode to episode. To accomplish the task, the agent must use observations of the environment to push boxes onto targets with its physical body. The level of difficulty of the task can be varied by changing the size and difficulty of the maze, the agent’s body, and the observations available to the agent.

In the base version of MuJoBan, the agent has a 2-DoF ball body that is able to move boxes by rolling into them (we use the ‘JumpingBallWithHead’ body described in [Tassa et al. \(2020\)](#) but with the jump action disabled). Boxes can move horizontally along the floor (including diagonally), but are constrained to not rotate. The agent can move boxes by pushing them with its body. Observations available to the agent include standard proprioceptive observations (see section B.4 for details), an egocentric camera (from the agent’s perspective), and a top-down camera that shows the entire level (see Figure 2). To allow the agent to infer its orientation from the egocentric

camera, the maze is surrounded by a fixed backdrop with a distinctive skyline pattern. In order to infer its orientation from the top-down camera view, the agent is topped with two differently colored ‘ears’ (a blue left ear and a red right ear), whose relative position indicates the agent’s current orientation. If the agent solves the puzzle by moving each box onto a pad, the agent receives a reward of 10 and the episode terminates. The agent also receives a shaping reward of 1 for pushing a box onto a pad or of -1 for pushing a box off of a pad. The agent receives 0 reward at all other transitions. The episode terminates if the level is unsolved after 45 seconds (900 control steps). While many easy levels can be solved in fewer than 10 seconds, more difficulty levels sometimes require close to the full 45 seconds to be completed.

Because the game runs in a physics engine, it is difficult or impossible to implement the rules of the abstract Sokoban perfectly. In this sense, MuJoBan relaxes the rules of Sokoban, and hence agents may ‘cheat’ by using moves that are impossible in Sokoban to solve the task in MuJoBan. For example, in some situations, an agent can ‘pull’ a stuck box out of a corner by rubbing against the box’s side and nudging it away from the wall with the resulting friction. In the base version of the game, we find that agents often solve the task by simultaneously pushing more than one box, by using diagonal movements to move between boxes, or by pushing boxes past each other while allowing them to partially occupy the same grid position in state space (some examples are shown in the video [here](#)).

To mitigate these problems, we also introduce a variant of the game in which physical pegs are inserted on the floor at every grid intersection point. Boxes cannot go through pegs, which makes diagonal motion effectively impossible. This forces the agent to carefully move boxes from one grid position to another in parallel to the axes of the grid. This game variant also makes it harder for the agent to push more than one box at once, and in practice we observed that agents solve this task using a strategy that more closely resembles a legal Sokoban strategy. Because of this, this game variant is also harder to solve.

The difficulty of MuJoBan can be varied in several other ways. First, the agent’s body can be easily swapped with any other available locomoting agent,<sup>1</sup> which can provide alternative, harder variants of MuJoBan, or tasks for transfer learning. Finally, while we test agents with access to all observations, variants could be made by restricting access to only a subset. In particular, not providing the top-down camera makes the problem more strongly partially observed, since the entire configuration of the maze (and the target configuration of the maze) cannot be inferred from a single first-person camera view. Level difficulty can be controlled by varying the size of the grid and the number of boxes per level. We procedurally generate levels as in ([Guez et al., 2019](#); [Racanière et al., 2017](#)). These reference levels fit within a  $10 \times 10$  grid, contain 4 boxes, and are of varying difficulty. Training with smaller, easier, levels is also supported, for example to create a curriculum.

## 2.2. MuJoXO

MuJoXO embeds the classic game of tic-tac-toe in a simulated physical environment. The agent plays by controlling a robotic arm to touch board positions on a tic-tac-toe board presented in front of the arm. Invisible touch pads are present around each location where a piece can validly be positioned on the board. The arm is based on a 9-DoF Jaco arm ([Campeau-Lecours et al., 2017](#)), controlled at 20Hz. When the arm makes contact with a touch pad at a position corresponding to a legal move, the move is registered and a piece (of the agent player’s color) appears at the pad’s physical location. Immediately following such an event, an opponent move is decided by an abstract planner and a piece is placed near the chosen location. The position of the piece on the game board is perturbed with random noise so that the pieces do not align perfectly on the grid (just as we would expect in a real-world realization of the game). After every move, and at the start of every episode, the arm is placed in a random configuration above the board.<sup>2</sup> To allow for different levels of difficulty, the epsilon-greedy opponent plays a random legal move with a configurable probability  $\epsilon$  at each step, and an optimal move otherwise.




Figure 3 | Example of the agent’s camera view in MuJoXO.

The inputs available to the agent are composed of proprioceptive sensory data (joint angles, velocities, and torques), the global coordinate of the end effector position, the abstract state of the board, as well as a camera

<sup>1</sup>Other bodies (‘walkers’) are available in the DeepMind control suite’s [locomotion task library](#). We provide example code in the appendix showing how to configure MuJoBan to use a different body.

<sup>2</sup>This simplifies the problem since it avoids the need to lift or untangle the arm after a move. This reset can easily be disabled.

facing the arm and the board which provides a RGB input (84x84 pixels by default). Although we provide all these inputs to the agent, some might not be strictly necessary and may simplify the task. A subset of them can be selected for increased realism and difficulty, but we found this version to be sufficiently challenging (see Section 5). Rewards are 0 for the duration of the episode, and a reward of 1 is given at the end of the game in case of a win or 0.5 in case of a draw. Episodes terminate after the game is decided or after 30 seconds. Games played with trained agents finish in  $\sim 5$  seconds, or  $\sim 100$  agent steps.

### 2.3. MuJoGo



Figure 4 | **Left:** example of the agent’s camera view in MuJoGo. **Right:** Top-down view of the board with all touch pads visualized. The ‘pass’ touch pads are shown as dark gray regions to the left and right of the board, and the touch pads for each intersection are shown as a checkerboard.

MuJoGo is set up similarly to tic-tac-toe, with the difference that the board size is a 7x7 Go board<sup>3</sup> and touches are registered around grid intersections. In addition, large touch pads on the left and right sides of the board are used to register pass moves. To encourage more consistent reaching movements across long Go games, only the end effectors can register a move (rather than any part of the arm, as in MuJoXO). Episodes terminate after 45 seconds during training, but we allow for up to 60 seconds per episode when evaluating agents. Timed out games are considered losses. An efficient agent can complete a game of MuJoXO in roughly 10 moves (about 0.5 seconds) or a game of MuJoGo in roughly 50 moves (about 2.5 seconds). Games are scored according to Tromp-Taylor rules (Tromp and Taylor, 1995) with a komi of 5.5. The opponent’s move selection is decided using the GNU Go program, a traditional Go program (Bump et al., 2009). Opponent strength is configured both with the GNU Go difficulty level  $L$  and the probability  $\epsilon$  of making a random legal action.<sup>4</sup> Captures are implemented by removing the appropriate stones from the board after each move.

To handle the underlying abstract game logic of tic-tac-toe and Go, we leverage the OpenSpiel framework (Lanctot et al., 2019). For our domains we chose to provide an opponent in the form of a configurable game engine rather than relying on self-play (as in Silver et al. 2017), to make our tasks more compatible with single-agent training frameworks. While we focus on tic-tac-toe and Go in this report, other board games could be straightforwardly included in this challenge suite. In their current form, both MuJoXO and MuJoGo employ certain simplifications: for instance, there are no collisions between board pieces (they are allowed to overlap on the board). Furthermore, picking and placing of pieces with the robotic hand is not required. Finally, the agent always plays the first move (black stones in Go). These simplifications could be removed in future versions of the tasks.

## 3. Background and related work

In this report, we describe a suite of domains designed to spur developments in long-horizon, embodied reasoning. While the components of this problem have been approached from many angles, little work in the RL and deep learning community has attempted to take on the problem as a whole. We believe that this is at least in part due to the paucity of suitable benchmarks. In this section, we review several recent approaches to motor control and long-horizon reasoning. This summary is not meant to be exhaustive, but rather to give a glimpse of the set of

<sup>3</sup>Here, we focus on the 7x7 board size, as we found it challenging for current algorithms. However, the code for this environment allows for larger board sizes, so the difficulty of this task can be easily scaled up in the future.

<sup>4</sup>Communication with GNU Go is made through the Go Text Protocol (GTP), other Go programs supporting GTP could easily be used instead.



approaches that may be important for mastering domains like MuJoBan, MuJoXO, and MuJoGo without using privileged information.

### 3.1. Challenges in motor control

The last few years have seen significant interest in approaches for learning control of simulated or real robots and other embodied systems. This includes, in particular, a significant body of work on robotic arm manipulation at different levels of complexity (e.g., [Lee et al., 2019a](#); [Popov et al., 2017](#); [Riedmiller et al., 2018](#); [Zhu et al., 2018](#)). Different publications consider control problems and solution strategies that emphasize different dimensions of the problem. These include a focus on dexterity (e.g., [Akkaya et al., 2019](#); [OpenAI et al., 2018](#)), but also long-horizon, multi-step control (e.g., [Duan et al., 2017](#); [Gupta et al., 2019](#); [Xu et al., 2017](#)). Compared to the environments presented in this report, the combinatorial complexity and horizon of most tasks considered in the literature is low. There is also a growing interest in multi-step, long-horizon control for locomotion and whole body control. Work such as ([Heess et al., 2017](#); [Peng et al., 2018](#)) considered agile but primarily short-horizon, reactive locomotion behavior. More recent work such as [Merel et al. \(2019, 2020\)](#) emphasize locomotion and whole-body control embedded in more complex visuomotor-reasoning tasks or multi-agent strategies (e.g. [Liu et al., 2019](#)).

Many of the problems studied in the motor control and robotics literature are challenging because of the difficulty of making conceptually simple, but high-level decisions in realistic state and action spaces, which may be continuous and high-dimensional. Consider a typical manipulation problem requiring a robot to pick up an object and place it in a box. This problem couples perception and fine-grained motor control. The underlying problem can in principle be solved with a simple procedure once the appropriate abstractions are identified (e.g. “move over the object, pick up the object, move over the box, drop the object”). However, non-trivial sequences of motor commands may be required to execute each abstraction. The problem is difficult largely because it requires identifying the underlying structure through perception and correctly executing a procedure with motor control.

### 3.2. Challenges in long-horizon reasoning

In contrast, in problems such as Go and chess a game is won or lost only after a long chain of decisions and actions, and the player may face a combinatorial number of game states that they have to respond to. For a plan to be successful, it typically needs to be executed with a high level of precision over intervals spanning a large number of decisions, as each decision can have a large impact on the final outcome, and the effects of a poorly made decision can be very difficult to recover from. These problems are notoriously difficult for RL algorithms because of the wide variety of situations that can be encountered, the difficulty of finding successful plays by chance, and the long delay between actions and rewards ([Arjona-Medina et al., 2019](#); [Hung et al., 2019](#); [Sutton and Barto, 2018](#)). This difficulty is exacerbated when planning problems are embedded in physical domains, as many small control steps may be required to carry out a single high-level decision. Waiting to reinforce the rare, random actions that lead to rewards may be theoretically possible, but it is unworkable in practice ([Hung et al., 2019](#); [Sutton and Barto, 2018](#)).

The robotics community has long recognized the importance of solving the joint problem of task-driven reasoning and movement/control (together known as *task and motion planning*; [Kambhampati et al. \(1991\)](#); [Latombe \(1991\)](#); [Wilfong \(1991\)](#); [Yap \(1986\)](#)). Researchers in robotics have studied these problems using reasoning domains that require long-horizon planning and multi-object sequential manipulation (e.g., [Dantam et al., 2018](#); [Lagriffoul et al., 2018](#)), and progress in this area has largely been driven by the design and engineering of domain-specific hand-crafted and structured methods ([Dantam et al., 2016](#); [Kaelbling and Lozano-Pérez, 2011](#); [Plaku and Hager, 2010](#); [Siméon et al., 2004](#); [Srivastava et al., 2014](#); [Toussaint, 2015](#)), which typically rely on access to privileged information. Although recent approaches to task and motion planning have incorporated methods from deep learning and RL (e.g., [Chitnis et al., 2016](#); [Driess et al., 2020](#); [Kim and Shimanuki, 2019](#); [Kim et al., 2019](#)), research on task and motion planning has largely occurred independently of developments in RL, and benchmarks from task and motion planning are typically not designed with RL methods in mind.

In the RL literature, a number of (often overlapping) lines of research consider strategies that may eventually address this challenge. For instance, research in model-based planning (e.g., [Chua et al., 2018](#); [Co-Reyes et al., 2018](#); [Finn and Levine, 2017](#); [Silver et al., 2016, 2017, 2018](#); [Sutton, 1991](#); [Wang and Ba, 2020](#)) studies how models can be used to identify good actions either to update the policy or to structure interactions with the environment. The use of models may improve data efficiency or facilitate transfer and may be particularly useful

in settings with combinatorial action choices and the need to make decisions that are not easily undone. Similarly, techniques developed in the hierarchical reinforcement learning literature (e.g., [Bacon et al., 2017](#); [Dayan and Hinton, 1993](#); [Florensa et al., 2017a](#); [Haarnoja et al., 2018](#); [Hausman et al., 2018](#); [Merel et al., 2019](#); [Nachum et al., 2019](#); [Sutton et al., 1999](#)) focus on structured behavior representations that may encourage learning by making behaviors more reusable, by providing structure for the credit assignment or exploration process, or by otherwise improving the abstractions learned by RL.

One major challenge when solving long-horizon reasoning problems is the so-called “exploration problem”, i.e. how to discover the rare strategies that lead to rewarded behaviors. While robust and directed exploration behavior can in principle arise solely by generalizing from relevant past experience, it is often difficult in practice. A range of methods have been proposed to bootstrap the learning process when training from scratch. One common solution is to use shaping rewards (e.g., [Popov et al., 2017](#)), i.e. rewards that are widely available in state space and which provide local information that guides the agent towards the solution, thereby shortening the effective horizon for credit assignment. Another common solution is to use a curriculum (e.g., [Florensa et al., 2017b](#); [Graves et al., 2017](#); [Riedmiller et al., 2018](#)), i.e. a set or sequence of tasks of increasing difficulty chosen such that mastering the simple tasks makes it easier for the agent to learn solutions to harder tasks, thus also guiding the agent to a solution of the target task. A similar role can be played by demonstrations of expert behaviour, which can be used to provide examples of trajectories that lead to reward or to initialize policies. This is a common approach ([Nair et al., 2018](#)), and has been used as a part of strategies for tackling StarCraft ([Vinyals et al., 2019](#)), Go ([Silver et al., 2016](#)), as well as motor control tasks ([Ho and Ermon, 2016](#); [Merel et al., 2020](#); [Peng et al., 2018](#)). Shaping rewards, curricula, and expert behavior explicitly or implicitly encode knowledge about the solution to a given task. An alternative line of work focuses on methods that drive the agent to search over the state-action in an effective manner (e.g., [Schmidhuber, 1991](#); [Thrun, 1992](#)). One popular way of doing so is to use intrinsic reward functions that incentivize the agent to visit different parts of the state space (e.g., [Bellemare et al., 2016](#); [Florensa et al., 2017a](#); [Gregor et al., 2017](#); [Hausman et al., 2018](#); [Strehl and Littman, 2008](#)).

### 3.3. Related RL benchmarks

The growing interest in motor control problems that require a combination of advanced motor skills, perception, and longer-horizon reasoning has led to the demand for suitable benchmark tasks. Whereas initial task suites focused on short-horizon single task challenges ([Brockman et al., 2016](#); [Tassa et al., 2018](#)), more recently released benchmark suites (e.g. [Fan et al., 2018](#); [James et al., 2020](#); [Yu et al., 2019](#)) focus on multi-task scenarios, where the same embodiment needs to solve a diverse set of challenges or is required to adapt existing skills to a new problem. Similar to our domains, the recently released IKEA furniture assembly benchmark ([Lee et al., 2019b](#)) emphasizes long-horizon reasoning, but does so in the context of an embodied construction task and provides a different trade-off in terms of perceptual difficulty, motor control, and the combinatorial complexity of the underlying planning problem. The DeepMind locomotion suite and soccer environment ([Liu et al., 2019](#); [Tassa et al., 2020](#)) include both short-horizon, agility-focused problems as well as longer-horizon and multi-agent challenges. Complementary to this task suite, which primarily focuses on motor control and multi-agent coordination, the tasks we introduce couple perception and physical interaction with a challenging abstract structure in a single-agent setting.

The above summary only provides a glimpse of the relevant body of work in the RL literature. With the environments described in section 2 we provide a test bed that we hope will allow for systematic comparisons of existing approaches and facilitate the development of new approaches, thus accelerating research into difficult embodied reasoning problems combining high-level reasoning, low-level motor control, and perception.

## 4. Experiments

The domains introduced in the previous section require the agent to learn a robust motor policy and align this policy with a high-level strategy that solves the puzzle (in the case of MuJoBan) or wins the game (in the case of MuJoXO and MuJoGo). The goal of this section is to establish baselines for these domains for future comparisons, and also as a way to better grasp the challenge behind these tasks. As we will show, naïve end-to-end training with state-of-the-art RL algorithms is largely unsuccessful, but an actor-critic architecture that incorporates expert planner information can learn good policies for each of these domains. This suggests that deep RL methods can learn appropriate motor plans when provided with shaping rewards that reflect high-level state and strategy

information.

To better understand the shortcomings of standard deep RL methods in these domains, we examine how much expert knowledge about the underlying abstract task (its state, dynamics, or solution) is needed to enable effective learning. To do so, we probe the performance of sensorimotor agents with access to three levels of information: (i) information about the abstract state, abstract dynamics, and the solution (the *expert planner* condition), (ii) information about the abstract state and dynamics, but not the solution (the *random planner* condition), (iii), minimal or no information about the abstract state (beyond the sensory observations) and no information about the dynamics or solution (the *vanilla agent* condition). We next describe the basic structure of our architecture (in section 4.1) before describing the details of the experimental conditions (in section 4.2).

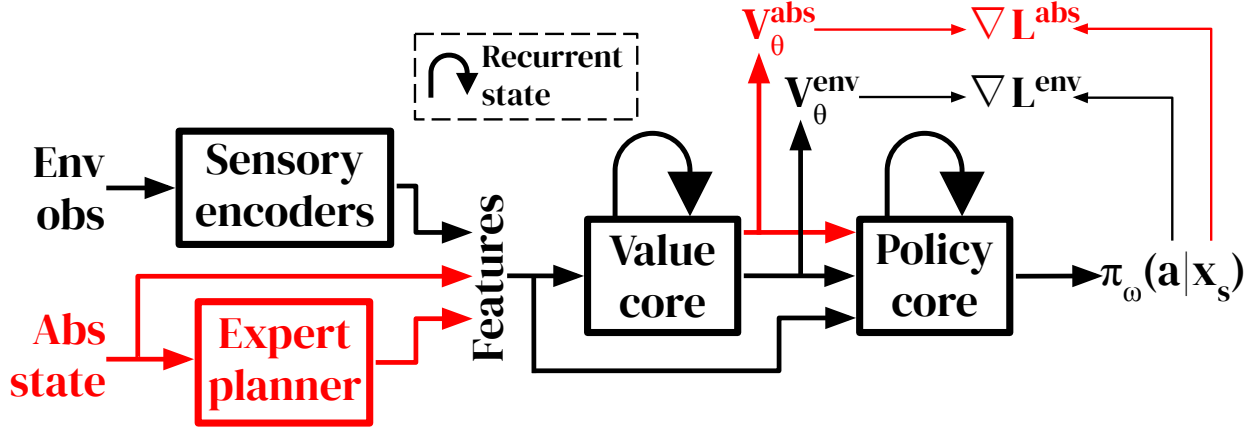


Figure 5 | The agent architecture used for the ‘expert planner’ condition. The standard components of the actor-critic policy gradient architecture are shown in **black**, while the components of the expert planner and the corresponding policy gradient computation are shown in **red**. The ‘vanilla agent’ condition uses only the components shown in black. The ‘random planner’ condition includes the components shown in red, but uses a random planner in place of the expert planner. Details of the architecture are given in Section B of the Appendix.

#### 4.1. Agent architecture

Our agent architecture is built around an actor-critic network structure similar to others used in the literature on RL for continuous control (e.g. Heess et al. (2017)). This architecture consists of several submodules designed for sensory encoding, value estimation, and continuous policy output. The full network architecture is trained end-to-end using the distributed IMPALA actor-critic algorithm with V-trace off-policy corrections (Espeholt et al., 2018).

There are two key differences between our architecture and those used elsewhere in continuous control. The first difference is the inclusion of an expert planner module that maps ground truth abstract states to target abstract states and the game actions needed to reach them. In the context of the continuous control task the agent needs to perform, the information fed into and returned from the expert should be viewed as *privileged* information. This is because the abstract game state and the action an expert would choose may be hard or impossible to infer from sensorimotor observations, but they are sufficient to solve the underlying abstract problem. We include this component to allow us to analyse the role that knowledge about the structure of the underlying state space and abstract actions plays in solving games that are grounded in a physical environment.

The second difference is the addition of an auxiliary task to follow the expert’s instructions in the abstract state space. In particular, the agent receives a reward of constant magnitude when the actions executed by the policy result in the state-space transition suggested by the expert. This reward is provided if the agent follows the suggestion within a fixed time limit. In addition to providing the agent with information about the solution to the task, this reward provides the agent with information about the structure of the abstract state space and the effect of low-level actions in this state space. We structure the auxiliary task using episodes that are shorter than those of the main task and which occur throughout the main episode (each auxiliary task episode can be thought of as a sub-episode of the main episode). The target of an auxiliary task is set by the planner at the beginning of each



auxiliary task episode. Auxiliary task episodes are reset when the agent reaches the target or when the time limit of the auxiliary episode is reached. Auxiliary task episodes are produced until the end of the main task episode is reached.

In early experiments, we estimated a single value that reflected the contribution of both the main task and the auxiliary task. We found that the agent struggled to learn in this setting. We believe this occurs because mixing the values for the two tasks makes it impossible to correctly bootstrap on either task and because the very different timescales of the two rewards makes it hard to balance their contributions to a single value. To encourage the agent to maximize both objectives, we estimate the value contributions of the two terms using two distinct network outputs (see Fig. 5) and optimize the two value estimates separately. The addition of a second value estimate more cleanly separates the value associated with following the planner’s suggestions from that associated with the task reward. It also allows us to handle discounting, bootstrapping, and episode termination for the two objectives separately. We found this to lead to more stable training. The agent architecture, including the two value estimates, is illustrated in Fig. 5.

The final agent update is produced by adding the policy gradient contributions of the two objectives (each of which uses the advantage estimates from the corresponding value output):

$$\nabla \mathcal{L}^{env} = \mathbb{E}_{x_s, a_s} [\rho_s \nabla_{\omega} \log \pi_{\omega}(a_s | x_s) (r_t^{env} + \gamma^{env} v_{s+1}^{env} - V_{\theta}^{env}(x_s))] \quad (1)$$

$$\nabla \mathcal{L}^{abs} = \mathbb{E}_{x_s, a_s} [\rho_s \nabla_{\omega} \log \pi_{\omega}(a_s | x_s) (r_t^{abs} + \gamma^{abs} v_{s+1}^{abs} - V_{\theta}^{abs}(x_s))] \quad (2)$$




Figure 6 | Illustration of different inputs in MuJoBan: **(left)** top-down camera, **(middle)** the corresponding abstract state, and **(right)** the target state provided by the expert planner. Color coding for the abstract state is as follows, grey: walls, blue: agent, yellow: boxes, red: box targets.

Eq. (1) describes the policy gradient for the main environmental task and eq. (2) describes the policy gradient for the abstract shaping task. We approximate the expectation over the agent’s full state ( $x_s$ ) and action ( $a_s$ ) occupancy using rollouts sampled from the agent’s experience using its current policy. We use different termination and discount factors  $\gamma$  for each gradient computation. Here,  $V_{\theta}(x_s)$  refers to the value approximation (with parameters  $\theta$ ) at state  $x_s$ ,  $v_{s+1}$  of the value target,  $r_t$  to the reward,  $a_s$  to the motor action,  $\pi_{\omega}$  to the policy (with parameters  $\omega$ ),  $\gamma$  to the discount factor, and  $\rho_s$  to the truncated importance sampling weight, as described by Espeholt et al. (2018).  $\nabla \mathcal{L}$  is the policy gradient (taken with respect to a proxy loss  $\mathcal{L}$  for each task, that corresponds to maximizing the return of that task). All variables in Eq. (1) with superscript ‘env’ are associated with the environment’s motor control task (solving the puzzle or winning the game) while those in Eq. (2) with superscript ‘abs’ are associated with the auxiliary task in the abstract state space.

In the expert-planner condition, the current abstract state and the target abstract state predicted by the expert planner (as well as the expert’s abstract action, in the case of MuJoBan) are given as input to the value and policy components of the architecture. The specified abstract states and actions reflect the underlying abstract task. For MuJoBan, this observation describes the full state of the Sokoban maze (including the position of the agent, the boxes, and the targets), while the action is the direction of the next movement (north, east, south, or west); the abstract inputs are illustrated in Figure 6. For MuJoGo and MuJoXO, the state describes the full board game, including the position of all pieces. For the board games, the expert does not return an explicit action, as the state is specified in global board coordinates and the action can be inferred directly from the current and next board configuration (this same property is exploited in the design of the agent architectures for board games in Silver

et al. (2016) and successive work). We use a pre-trained Sokoban RL agent (with the convLSTM-based architecture described in Guez et al. (2019)) as the expert planner for MuJoBan and we use open-source, off-the-shelf agents for the board games: GNU Go (level 10) (Bump et al., 2009) for MuJoGo and a minimax tree-search agent for MuJoXO. The expert planner is only used to provide target states and actions during agent training: we do not train it or differentiate through it.

The agent receives sensory input from proprioceptive sensors (including touch, position, velocity, and acceleration sensors) as well as cameras (ego-centric and top-down cameras in the case of MuJoBan). Each of these inputs are processed via a multi-layer Perceptron (for non-visual sensory inputs) or a residual network (He et al., 2016) (for visual inputs) and concatenated before being passed to the recurrent cores. We give details of the observations and architecture in Section B of the Appendix.

## 4.2. Experimental conditions

In the experiments, we probe the effect of different types of expert information on task performance. We modify the architecture as follows in each of the three conditions:

- **Expert planner agent:** the agent has access to an expert planner that consistently solves the abstract task embedded in the environment. This agent uses the full architecture described in Section 4.1 and shown in Fig. 5, including the reward (and corresponding policy gradient term) reflecting how well the agent follows the expert plan in the abstract space.
- **Random planner agent:** the agent receives the current abstract state and a target abstract state chosen from among the possible valid next states. The agent architecture and reward in this case is identical to the previous case; the only difference is that target states are chosen randomly instead of by an expert. Although these random high-level actions sometimes conflict with the goals of the underlying game, we reason that because this extra reward provides information about the underlying game and its dynamics, it may help the agent explore the state space in more structured manner.
- **Vanilla agent:** the agent has access to minimal prior knowledge about the underlying game and no access to the planner. In the case of MuJoBan, the agent is not given the abstract game state. In the case of MuJoXO and MuJoGo, the agent is given the current board state to simplify learning, though we expect that similar results could be obtained using only camera inputs. This results in a standard recurrent actor-critic architecture with minimal prior knowledge about the abstract game and no direct information about the dynamics of the game. We train this architecture to maximize the environment reward using only a single value estimate.

In MuJoBan, we train agents with a curriculum composed of mazes with 5 difficulty levels, cf. Appendix Table 1. During evaluation, we track the agent’s performance on each of the difficulty levels separately. In MuJoGo, we set the opponent parameters to  $\epsilon = 0$  and  $L = 10$  for the GnuGo opponent. In MuJoXO, the opponent plays random moves with probability  $\epsilon = 0.25$  and plays optimal tic-tac-toe moves otherwise.

## 5. Results

We report our experimental results below, organized by domain.

### 5.1. MuJoBan

Our experimental results for MuJoBan confirm that the task is challenging for current RL methods. Learning curves for the larger grids are shown in Figure 7, and for other grid sizes in Appendix Figure 9. In the grid peg version, which forces a closer correspondence to the abstract Sokoban problem, agents were unable to solve more than half of the levels, even after extensive training and even though these levels correspond to relatively straightforward abstract problems. Vanilla RL agents, which do not have any abstract augmentation, were inefficient at learning both versions of the task (with and without pegs). This was true even in the presence of a curriculum designed to make it easier for the agent to learn the important skills needed (e.g., how to push boxes to targets). The curriculum provides the agent with a signal that is sufficiently strong to learn the basic mechanics of the task, which suggests that the agent’s failure to solve the harder levels is not primarily due to limited exploration. Rather, this suggests

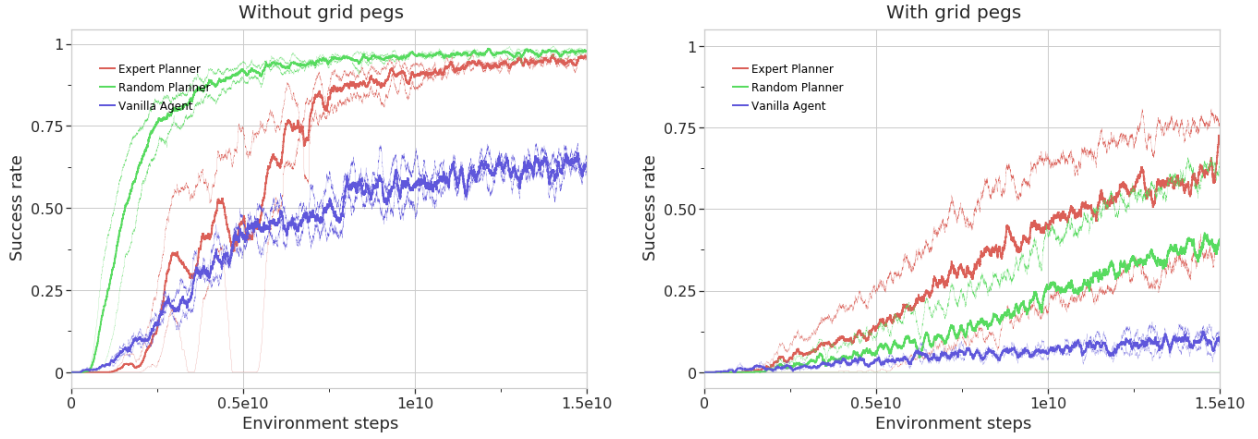


Figure 7 | Agent performance on MuJoBan with the larger 10x10 levels (difficulty 5). In the peg-less version (a), the vanilla agent only learned to solve half of the levels after extensive training. The random planner agent performed best in this setting. In (b), where the pegs constrain the agent to solve the task with logical constraints closer to the abstract game, the expert planner performed better and the vanilla planner only ever learned to solve a small fraction of the tasks. Thin traces show the results of each individual seed and thick traces show the mean results across all seeds of each condition. Due to the long wall-clock times associated with running these experiments to convergence, we terminated each run after  $1.5 \times 10^{10}$  steps (approx. 11 days using 2x2 TPU v2 slices).

that the kind of generalization needed to move beyond greedy behavior (often sufficient to solve simple levels) to longer-term reasoning (needed to solve harder levels) does not emerge in our model-free MuJoBan agent. For context, model-free agents can obtain 90%+ success on the hardest levels in the original, non embodied version of Sokoban (Guez et al., 2019). Moreover, the control problem associated with implementing the individual steps of an optimal policy in the abstract task is not particularly difficult (Karkus et al. (2020) provides some evidence for this assertion). It appears that the RL agent’s difficulty in mastering this task arises due to the combination of multi-step reasoning and control, rather than either of these in isolation.

The version without grid pegs was markedly less demanding for all agents, presumably because there is more flexibility to solve the level in an unconventional way (e.g., push boxes diagonally, recover from dead end scenarios). Nevertheless, the model-free agent did not progress much beyond the 50% mark whereas privileged agents (as described in Section 4) achieved good performance (around 90% of level solved). In the peg-less version, we also found the random planner to perform more robustly than the expert planner — perhaps because this version aligns less directly with the abstract version of the problem that the expert planner guides towards. In the version with pegs, the agent equipped with the expert planner learned faster – perhaps because the expert subgoals are often on the optimal solution path in this physics version of the problem. We summarize the final performance on all conditions in Appendix Table 4.

## 5.2. MuJoXO

MuJoXO has a modest planning load compared to MuJoGo and MuJoBan, since the abstract problem space of tic-tac-toe is relatively small (with an average branching factor of 4 and depth less than 9). Nevertheless, careful decision making is still needed given the presence of irreversible decisions and the added requirement of placing moves with a robotic arm. Indeed, we observed that our model-free agent required a large number of games (on the order of 1 million) before it could play convincingly, even in this limited planning setting. This is shown in Figure 8-a. In contrast, the agent equipped with an optimal abstract planner achieved that performance level with around 5 times less games of experience. The agent suggesting random abstract goals, corresponding to placing a legal move at random, was similarly helpful in early training, after which the agent plateaued without progress. Unlike MuJoBan, where physical states have a somewhat fuzzier correspondence to abstract states due to mismatches between the abstract and physical games, the subgoals suggested by the expert in MuJoXO are exactly aligned with the intermediate states of an optimal strategy: the only way to win is to have placed the stone in the right square, and the exact physical position of that stone has then no bearing on the rest of the episode.




Figure 8 | Left: Agent performance on MuJoXO. Right: Agent performance on MuJoGo (random planner and vanilla agent curves are near 0). In both panels, thin traces show the results of each individual seed and thick traces show the mean results across all seeds of each condition. Please note the difference in  $y$  axis scale for each panel.

Our results suggest that knowledge of the abstract structure of the game and its dynamics aids learning, but is not sufficient to encourage the agent to discover how to reliably beat its opponent.

### 5.3. MuJoGo

In MuJoGo, the additional complexity of planning and the longer game times, considerably increased the learning difficulty for all agents. The vanilla (and random planner) agents did not show any sign of progress after billions of steps of training. Since the agent plays against a fixed-strength opponent in MuJoGo, exploration is a notable issue early on in training: the agent needs to observe wins in order to correctly adapt its actions. For the vanilla model-free agent, wins happened very infrequently (approximately 1 in every 1000 games), which might be sufficient to explain the poor performance we observed. The agent equipped with an abstract planner fared considerably better, eventually beating the fixed opponent in this 7x7 Go game, but it is still remarkably inefficient at learning, considering the small size of the board and the limited strength of its opponent. It required around 4M games to reach a win rate of  $\sim 60\%$  (the level of performance at the end of training reported in Figure 8). In comparison, 4M games is approximately the number of self-play games used by AlphaGo Zero to reach professional level on 19x19 boards in the abstract version of the problem (Silver et al., 2017).

## 6. Conclusion

Advances in deep learning have spurred profound progress in the field of reinforcement learning. The use of powerful function approximators in the form of large and deep neural networks has allowed machine learning practitioners to move away from hand-crafted features and to solve RL problems that require complex generalization for perception and control. However, many of these successes are limited to problems where decisions and control take place over relatively short time scales, and actions are taken in an abstracted state and action space tailored to the task. Problems that require reasoning and decision making over long time scales using sensorimotor control cannot yet be solved in an end-to-end fashion. These problems arise frequently in human behavior but are still hard to frame and rarely studied in a controlled experimental setting.

We believe that board games, which are among the most widely used test beds for studying planning capabilities in AI (Yannakakis and Togelius, 2018), are also well-suited for study in the embodied control context. While games like these are usually studied in an abstract form, they typically originate from real-world games that employ physical objects with real pieces played using human actuation. These physical version of a board game naturally capture the challenges of coupling abstract reasoning with motor control. They simultaneously provide well defined problems that are known to benefit from algorithmic advances. Because the abstract structure underlying many

board games is well understood, board games allow the researcher to easily tune the difficulty of the reasoning task and to analyze the strategies learned by agents.

To encourage research in this direction, we have introduced a set of board games (tic-tac-toe, Go, and Sokoban) embedded in physical simulations. We have demonstrated that a standard deep RL algorithm struggles to solve these games when they are physically embedded and temporally extended. Agents given access to explicit expert information about the underlying state and action spaces can learn to play these games successfully, albeit after extensive training. We present this gap as a challenge to RL researchers: What new technique or combination of techniques are missing from the standard RL toolbox that will close this gap? We open source these environments to spur development of agents that can reach high level of performance, without privileged information.

Humans appear to use a variety of approaches to help them reason about and solve physically embedded and temporally extended planning problems. They often receive hints about important sub-goals, directly from others in the form of demonstrations, or through culturally-specified cues such as the reuse of conventions in game boards, or abstract rule specifications via language. Humans also appear to discover potentially interesting structure through intrinsically motivated exploration and curiosity. By exploring and exploiting this structure, humans are able to observe, reason, and control in a more directed fashion. We hope that the environments provided here will spur research into how to coherently introduce these capacities into the next generation of RL agents.

## 7. Acknowledgments

We thank Yuval Tassa and Animesh Garg for helpful feedback on earlier versions of the paper.

## References

- I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, et al. Solving Rubik’s Cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.
- J. A. Arjona-Medina, M. Gillhofer, M. Widrich, T. Unterthiner, J. Brandstetter, and S. Hochreiter. RUDDER: Return decomposition for delayed rewards. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2019.
- P.-L. Bacon, J. Harb, and D. Precup. The option-critic architecture. In *Proceedings of AAAI Conference on Artificial Intelligence*, 2017.
- C. Beattie, J. Z. Leibo, D. Teplyaev, T. Ward, M. Wainwright, H. Küttler, A. Lefrancq, S. Green, V. Valdés, A. Sadik, J. Schrittwieser, K. Anderson, S. York, M. Cant, A. Cain, A. Bolton, S. Gaffney, H. King, D. Hassabis, S. Legg, and S. Petersen. DeepMind Lab. *arXiv preprint arXiv:1612.03801*, 2016.
- M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos. Unifying count-based exploration and intrinsic motivation. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2016.
- G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. OpenAI Gym. *arXiv preprint arXiv:1606.01540*, 2016.
- R. A. Brooks. Elephants don’t play chess. *Robotics and Autonomous Systems*, 6:3–15, 1990.
- D. Bump, G. Farneback, A. Bayer, and more. GNUGo 3.8. <https://www.gnu.org/software/gnugo/>, 2009.
- A. Campeau-Lecours, H. Lamontagne, S. Latour, P. Fautoux, V. Maheu, F. Boucher, C. Deguire, and L.-J. C. L’Ecuyer. Kinova modular robot arms for service robotics applications. *International Journal of Robotics Applications and Technologies (IJRAT)*, 5(2):49–71, 2017.
- R. Chitnis, D. Hadfield-Menell, A. Gupta, S. Srivastava, E. Groshev, C. Lin, and P. Abbeel. Guided search for task and motion plans using learned heuristics. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2016.
- K. Chua, R. Calandra, R. McAllister, and S. Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2018.
- J. D. Co-Reyes, Y. Liu, A. Gupta, B. Eysenbach, P. Abbeel, and S. Levine. Self-consistent trajectory autoencoder: Hierarchical reinforcement learning with trajectory embeddings. In *Proceedings of International Conference on Machine Learning (ICML)*, 2018.



- N. T. Dantam, P. Kolhe, and M. Stilman. The Motion Grammar for physical human-robot games. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2011.
- N. T. Dantam, Z. K. Kingston, S. Chaudhuri, and L. E. Kavraki. Incremental task and motion planning: A constraint-based approach. In *Proceedings of Robotics: Science and Systems*, 2016.
- N. T. Dantam, S. Chaudhuri, and L. E. Kavraki. The task-motion kit: An open source, general-purpose task and motion-planning framework. *IEEE Robotics & Automation Magazine*, 25(3):61–70, 2018.
- P. Dayan and G. E. Hinton. Feudal reinforcement learning. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 1993.
- D. Driess, H. Jung-Su, and M. Toussaint. Deep visual reasoning: Learning to predict action sequences for task and motion planning from an initial scene image. In *Proceedings of Robotics: Science and Systems*, 2020.
- Y. Duan, M. Andrychowicz, B. C. Stadie, J. Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba. One-shot imitation learning. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2017.
- L. Espeholt, H. Soyer, R. Munos, K. Simonyan, V. Mnih, T. Ward, Y. Doron, V. Firoiu, T. Harley, I. Dunning, S. Legg, and K. Kavukcuoglu. IMPALA: scalable distributed deep-RL with importance weighted actor-learner architectures. In *Proceedings of International Conference on Machine Learning (ICML)*, 2018.
- L. Fan, Y. Zhu, J. Zhu, Z. Liu, O. Zeng, A. Gupta, J. Creus-Costa, S. Savarese, and L. Fei-Fei. SURREAL: Open-source reinforcement learning framework and robot manipulation benchmark. In *Conference on Robotic Learning (CoRL)*, 2018.
- N. Fazeli, M. Oller, J. Wu, Z. Wu, J. Tenenbaum, and A. Rodriguez. See, feel, act: Hierarchical learning for complex manipulation skills with multisensory fusion. *Science Robotics*, 4(26):eaav3123, 2019.
- C. Finn and S. Levine. Deep visual foresight for planning robotic motion. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2017.
- C. Florensa, Y. Duan, and P. Abbeel. Stochastic neural networks for hierarchical reinforcement learning. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2017a.
- C. Florensa, D. Held, M. Wulfmeier, M. Zhang, and P. Abbeel. Reverse curriculum generation for reinforcement learning. In *Conference on Robotic Learning (CoRL)*, 2017b.
- A. Graves, M. G. Bellemare, J. Menick, R. Munos, and K. Kavukcuoglu. Automated curriculum learning for neural networks. In *Proceedings of International Conference on Machine Learning (ICML)*, 2017.
- K. Gregor, D. J. Rezende, and D. Wierstra. Variational intrinsic control. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2017.
- K. Gregor, D. J. Rezende, F. Besse, Y. Wu, H. Merzic, and A. van den Oord. Shaping belief states with generative environment models for RL. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2019.
- A. Guez, M. Mirza, K. Gregor, R. Kabra, S. Racanière, T. Weber, D. Raposo, A. Santoro, L. Orseau, T. Eccles, G. Wayne, D. Silver, and T. P. Lillicrap. An investigation of model-free planning. In *Proceedings of International Conference on Machine Learning (ICML)*, 2019.
- A. Gupta, V. Kumar, C. Lynch, S. Levine, and K. Hausman. Relay policy learning: Solving long-horizon tasks via imitation and reinforcement learning. In *Conference on Robotic Learning (CoRL)*, 2019.
- T. Haarnoja, K. Hartikainen, P. Abbeel, and S. Levine. Latent space policies for hierarchical reinforcement learning. In *Proceedings of International Conference on Machine Learning (ICML)*, 2018.
- K. Hausman, J. T. Springenberg, Z. Wang, N. Heess, and M. Riedmiller. Learning an embedding space for transferable robot skills. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2018.
- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- N. Heess, D. TB, S. Sriram, J. Lemmon, J. Merel, G. Wayne, Y. Tassa, T. Erez, Z. Wang, S. M. A. Eslami, M. A. Riedmiller, and D. Silver. Emergence of locomotion behaviours in rich environments. *arXiv preprint arXiv:1707.02286*, 2017.
- J. Ho and S. Ermon. Generative adversarial imitation learning. In *Proceedings of Neural Information Processing*

*Systems (NeurIPS)*, 2016.

S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

C.-C. Hung, T. Lillicrap, J. Abramson, Y. Wu, M. Mirza, F. Carnevale, A. Ahuja, and G. Wayne. Optimizing agent behavior over long time scales by transporting value. *Nature communications*, 10(5223), 2019.

J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26), 2019.

S. James, Z. Ma, D. R. Arrojo, and A. J. Davison. RL Bench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters*, 2020.

A. Junghanns and J. Schaeffer. Sokoban: A challenging single-agent search problem. In *IJCAI Workshop on Using Games as an Experimental Testbed for AI Research*, 1997.

L. P. Kaelbling and T. Lozano-Pérez. Hierarchical task and motion planning in the now. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2011.

S. Kambhampati, M. R. Cutkosky, M. Tenenbaum, and S. Lee. Combining specialized reasoners and general purpose planners: a case study. In *Proceedings of AAAI Conference on Artificial Intelligence*, 1991.

P. Karkus, M. Mirza, A. Guez, A. Jaegle, T. Lillicrap, L. Buesing, N. Heess, and T. Weber. Deep modular reinforcement learning for physically embedded abstract reasoning. In *ICLR Workshop on Beyond Tabula Rasa in RL*, 2020. <http://www.betr-rl.ml/assets/camera-ready/19.pdf>.

B. Kim and L. Shimanuki. Learning value functions with relational state representations for guiding task-and-motion planning. In *Conference on Robotic Learning (CoRL)*, 2019.

B. Kim, L. P. Kaelbling, and T. Lozano-Pérez. Adversarial actor-critic method for task and motion planning problems using planning experience. In *Proceedings of AAAI Conference on Artificial Intelligence*, 2019.

F. Lagriffoul, N. T. Dantam, C. Garrett, A. Akbari, S. Srivastava, and L. E. Kavraki. Platform-independent benchmarks for task and motion planning. *IEEE Robotics and Automation Letters*, 3(4):3765–3772, 2018.

M. Lanctot, E. Lockhart, J.-B. Lespiau, V. Zambaldi, S. Upadhyay, J. Pérolat, S. Srinivasan, F. Timbers, K. Tuyls, S. Omidshafiei, et al. OpenSpiel: A framework for reinforcement learning in games. *arXiv preprint arXiv:1908.09453*, 2019.

J.-C. Latombe. *Robot Motion Planning*. Springer, 1991.

Y. Lee, E. S. Hu, Z. Yang, and J. J. Lim. To follow or not to follow: Selective imitation learning from observations. In *Conference on Robotic Learning (CoRL)*, 2019a.

Y. Lee, E. S. Hu, Z. Yang, A. Yin, and J. J. Lim. IKEA furniture assembly environment for long-horizon complex manipulation tasks. *arXiv preprint arXiv:1911.07246*, 2019b.

S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, (17):1–40, 2016.

S. Liu, G. Lever, J. Merel, S. Tunyasuvunakool, N. Heess, and T. Graepel. Emergent coordination through competition. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2019.

J. Merel, A. Ahuja, V. Pham, S. Tunyasuvunakool, S. Liu, D. Tirumala, N. Heess, and G. Wayne. Hierarchical visuomotor control of humanoids. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2019.

J. Merel, S. Tunyasuvunakool, A. Ahuja, Y. Tassa, L. Hasenclever, V. Pham, T. Erez, G. Wayne, and N. Heess. Reusable neural skill embeddings for vision-guided whole body movement and object manipulation. *ACM Transactions on Graphics*, 2020.

O. Nachum, S. Gu, H. Lee, and S. Levine. Near-optimal representation learning for hierarchical reinforcement learning. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2019.

A. Nair, B. McGrew, M. Andrychowicz, W. Zaremba, and P. Abbeel. Overcoming exploration in reinforcement learning with demonstrations. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2018.

OpenAI, M. Andrychowicz, B. Baker, M. Chociej, R. Józefowicz, B. McGrew, J. W. Pachocki, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder, L. Weng, and W. Zaremba. Learning

dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2018.

T. L. Paine, B. S. Caglar Gulcehre, M. Denil, M. Hoffman, H. Soyer, R. Tanburn, S. Kapturowski, N. Rabinowitz, D. Williams, G. Barth-Maron, Z. Wang, N. de Freitas, and W. Team. Making efficient use of demonstrations to solve hard exploration problems. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2020.

X. B. Peng, P. Abbeel, S. Levine, and M. van de Panne. DeepMimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics*, 37(4):143:1–143:14, 2018.

E. Plaku and G. D. Hager. Sampling-based motion and symbolic action planning with geometric and differential constraints. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2010.

I. Popov, N. Heess, T. P. Lillicrap, R. Hafner, G. Barth-Maron, M. Vecerík, T. Lampe, Y. Tassa, T. Erez, and M. A. Riedmiller. Data-efficient deep reinforcement learning for dexterous manipulation. *arXiv preprint arXiv:1704.03073*, 2017.

B. Quack, F. Wörgötter, and A. Agostini. Simultaneously learning at different levels of abstraction. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015.

S. Racanière, T. Weber, D. Reichert, L. Buesing, A. Guez, D. J. Rezende, A. P. Badia, O. Vinyals, N. Heess, Y. Li, et al. Imagination-augmented agents for deep reinforcement learning. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2017.

M. A. Riedmiller, R. Hafner, T. Lampe, M. Neunert, J. Degraeve, T. V. de Wiele, V. Mnih, N. Heess, and J. T. Springenberg. Learning by playing - solving sparse reward tasks from scratch. In *Proceedings of International Conference on Machine Learning (ICML)*, 2018.

J. Schmidhuber. Curious model-building control systems. In *Proceedings of International Joint Conference on Neural Networks*, 1991.

D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. v. d. Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, H. Fan, L. Sifre, G. v. d. Driessche, T. Graepel, and D. Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, 2017.

D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419):1140–1144, 2018.

T. Siméon, J.-P. Laumond, J. Corés, and A. Sahbani. Manipulation planning with probabilistic roadmaps. *The International Journal of Robotics Research*, 23(7-8):729–746, 2004.

H. F. Song, A. Abdolmaleki, J. T. Springenberg, A. Clark, H. Soyer, J. W. Rae, S. Noury, A. Ahuja, S. Liu, D. Tirumala, N. Heess, D. Belov, M. Riedmiller, and M. M. Botvinick. V-MPO: On-policy maximum a posteriori policy optimization for discrete and continuous control. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2020.

S. Srivastava, E. Fang, L. Riano, R. Chitnis, S. Russell, and P. Abbeel. Combined task and motion planning through an extensible planner-independent interface layer. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2014.

A. L. Strehl and M. L. Littman. An analysis of model-based interval estimation for markov decision processes. *Journal of Computer and System Sciences*, 74(8):1309–1331, 2008.

R. S. Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART Bulletin*, 2(4):160–163, 1991.

R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT Press, 2nd edition, 2018.

R. S. Sutton, D. Precup, and S. Singh. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1):181–211, 1999.

Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. de Las Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq,

- T. P. Lillicrap, and M. A. Riedmiller. DeepMind control suite. *arXiv preprint arXiv:1801.00690*, 2018. URL [https://github.com/deepmind/dm\\_control](https://github.com/deepmind/dm_control).
- Y. Tassa, S. Tunyasuvunakool, A. Muldal, Y. Doron, S. Liu, S. Bohez, J. Merel, T. Erez, T. Lillicrap, and N. Heess. dm\_control: Software and tasks for continuous control. *arXiv preprint arXiv:2006.12983*, 2020. URL [https://github.com/deepmind/dm\\_control](https://github.com/deepmind/dm_control).
- S. B. Thrun. Efficient exploration in reinforcement learning. *CMU Technical Report: CMU-CS-92-102*, 1992.
- E. Todorov, T. Erez, and Y. Tassa. MuJoCo: A physics engine for model-based control. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012.
- M. Toussaint. Logic-geometric programming: an optimization-based approach to combined task and motion planning. In *Proceedings of International Joint Conference on Artificial Intelligence*, 2015.
- J. Tromp and B. Taylor. Tromp-Taylor Go rules. <https://senseis.xmp.net/?TrompTaylorRules>, 1995.
- S. Tunyasuvunakool. LabMaze. <https://github.com/deepmind/labmaze>, 2019.
- O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, et al. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782): 350–354, 2019.
- T. Wang and J. Ba. Exploring model-based planning with policy networks. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2020.
- G. Wayne, C. Hung, D. Amos, M. Mirza, A. Ahuja, A. Grabska-Barwinska, J. W. Rae, P. Mirowski, J. Z. Leibo, A. Santoro, M. Gemici, M. Reynolds, T. Harley, J. Abramson, S. Mohamed, D. J. Rezende, D. Saxton, A. Cain, C. Hillier, D. Silver, K. Kavukcuoglu, M. Botvinick, D. Hassabis, and T. P. Lillicrap. Unsupervised predictive memory in a goal-directed agent. *arXiv preprint arXiv:1803.10760*, 2018.
- G. Wilfong. Motion planning in the presence of movable obstacles. *Annals of Mathematics and Artificial Intelligence*, (3):131–150, 1991.
- D. Xu, S. Nair, Y. Zhu, J. Gao, A. Garg, L. Fei-Fei, and S. Savarese. Neural task programming: Learning to generalize across hierarchical tasks. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2017.
- G. N. Yannakakis and J. Togelius. *Artificial Intelligence and Games*. Springer Publishing Company, Incorporated, 1st edition, 2018.
- C.-K. Yap. Algorithmic motion planning. In J. T. Schwartz and C.-K. Yap, editors, *Algorithmic and Geometric Aspects of Robotics*. Routledge, 1986.
- T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine. Meta-World: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robotic Learning (CoRL)*, 2019.
- Y. Zhu, Z. Wang, J. Merel, A. Rusu, T. Erez, S. Cabi, S. Tunyasuvunakool, J. Kramár, R. Hadsell, N. de Freitas, et al. Reinforcement and imitation learning for diverse visuomotor skills. In *Proceedings of Robotics: Science and Systems*, 2018.

## A. Code and Environment details

The code for all environments is available at [https://github.com/deepmind/dm\\_control](https://github.com/deepmind/dm_control). Below we provide examples and details for each environment. Each code example relies on the libraries included in the `dm_control` package (Tassa et al., 2020), including PyMJCF and Composer.

### A.1. MuJoBan

The code snippet below shows how to instantiate a MuJoBan environment. MuJoBan is a composer task. As arguments it requires a walker and maze. Here, the maze is constructed using one of the previously released BoxoBan levels (<https://github.com/deepmind/boxoban-levels>). And the walker provided is jumping ball walker from `dm_control.locomotion`.

```
from dm_control import composer
from dm_control.locomotion import walkers
from physics_planning_games.mujoban.mujoban import Mujoban
from physics_planning_games.mujoban.mujoban_level import MujobanLevel
from physics_planning_games.mujoban.boxoban import boxoban_level_generator

walker = walkers.JumpingBallWithHead(add_ears=True, camera_height=0.25)
maze = MujobanLevel(boxoban_level_generator)
task = Mujoban(walker=walker,
               maze=maze,
               control_timestep=0.1,
               top_camera_height=96,
               top_camera_width=96)
env = composer.Environment(time_limit=1000, task=task)
```

MuJoBan makes use of the LabMaze utility (Beattie et al., 2016; Tunyasuvunakool, 2019) to create textured 3D mazes from simple ASCII descriptions. Wall and floor textures of different colors are sampled randomly at the start of every episode.

### A.2. MuJoXO and MuJoGo

```
from physics_planning_games import board_games

environment_name = 'go_7x7'
env = board_games.load(environment_name=environment_name)
```

The simulated Jaco arm (Campeau-Lecours et al., 2017) used within these environments is part of the `dm_control` suite (Tassa et al., 2020).

## B. Experimental details

Table 1 | MuJoBan difficulty level settings. All levels smaller than 10×10 are padded with walls.

Level difficulty	Grid size	Number of boxes	Training ratio
1	5×5	1	0.25
2	7×7	1	0.25
3	7×7	2	0.2
4	8×8	3	0.2
5	10×10	4	0.1

### B.1. State estimation in MuJoBan

We estimate the abstract state from the physical scene in order to provide the symbolic input to the expert planner. We use the true physical location of the objects and round them off to their nearest grid positions. To prevent overlaps, we first estimate the boxes’ positions and then the player location. If an overlap occurs (around grid boundaries), we keep the old player’s abstract position.



## B.2. Expert Planner in MuJoBan

The agent has a time limit of  $T^{abs}$  to reach the given target state. When the agent reaches the target state or  $T^{abs}$  has been reached, we do another state estimation and a new target is proposed by expert planner. We keep the abstract state and target state constant within the period of length  $T^{abs}$ . The value we found to work best was  $T^{abs} = 50$ .

## B.3. Learning setup

We train in a distributed RL setup using 1000 actors generating trajectories chunked in unrolls of length 100. Parameters are updated in the learner using the RMSProp optimizer accumulated over batch sizes 64 (MuJoXO and MuJoGo) or 128 (MuJoBan).

## B.4. Other experimental details

Task time limits are 900 control steps for MuJoBan and MuJoGo, and 600 for MuJoXO. Each control timestep is 0.05 seconds for all domains. The discount factors used for all domains and auxiliary tasks are listed in Table 2.

All cameras (top-down, egocentric, or front-facing) provide RGB inputs with resolution 96x96 to the agents. For MuJoXO and MuJoGo, the proprioceptive inputs are the arm and hand joints' torque, position and velocity, as well of the position of the pinch site. For MuJoBan, the proprioceptive inputs are the joints' torque, velocity and position, as well as a touch sensor, accelerometer, gyroscope, and velocimeter.

In all experiments, we used network architectures that have been successfully used for other continuous control tasks using vision (Merel et al., 2019, 2020). **Sensory encoders:** we process all input images using a simplified residual network (He et al., 2016), whose parameters are given in Table 3. All proprioceptive inputs are concatenated, passed through an arcsinh nonlinearity to mitigate the effect of large magnitude values, and then processed by a single linear layer with a ReLU activation. When used as an input, the planner state is processed by a single convolutional layer with a 4x4 kernel and 100 units and a ReLU activation function. The output of this layer is flattened and passed through a linear layer with 256 units and a ReLU activation function. The output of all sensory encoders, as well as the state and expert planner output, are concatenated before being passed to the value and policy cores. **Value and policy architecture:** the value core is a long-short term memory (LSTM) recurrent neural network (Hochreiter and Schmidhuber, 1997) with 256 units and the policy core is an LSTM with 128 units. The output of the policy core parameterizes the mean and variance of a Gaussian distribution. Before being used to parameterize the Gaussian, the mean produced by the network is passed through a hyperbolic tangent nonlinearity, and the variance produced by the network is passed through a sigmoid nonlinearity and scaled to [0.1, 1].

We tuned the scale of the expert reward, the learning rate, and the two discount factors separately for each task using ad hoc grid sweeps. The experiments reported in here were run using the best combination of parameters we discovered using at least two random seeds (for MuJoBan, MuJoXO, and MuJoGo).

Table 2 | Discount parameters for each domain used in the experiment of Section 4.

Domain	$\gamma^{env}$	$\gamma^{abs}$
MuJoBan	0.99	0.9
MuJoXO	0.99	0.98
MuJoGo	0.99	0.98

### C. Additional figures

Table 3 | The residual network architecture used to process images in all experiments. The image resolution is reduced by half at each residual block group. Residual connections are made by adding the input of each residual block (in square brackets) to its output. Unlike other resnets, this architecture does not include any normalization layers.

input size	layer group
96×96	conv(3×3, 16) max_pool(3×3, stride 2) $\left[ \begin{array}{c} \text{ReLU} \\ \text{conv}(3\times 3, 16) \\ \text{ReLU} \\ \text{conv}(3\times 3, 16) \end{array} \right] \times 2$
48×48	conv(3×3, 32) max_pool(3×3, stride 2) $\left[ \begin{array}{c} \text{ReLU} \\ \text{conv}(3\times 3, 32) \\ \text{ReLU} \\ \text{conv}(3\times 3, 32) \end{array} \right] \times 2$
24×24	conv(3×3, 32) max_pool(3×3, stride 2) $\left[ \begin{array}{c} \text{ReLU} \\ \text{conv}(3\times 3, 32) \\ \text{ReLU} \\ \text{conv}(3\times 3, 32) \end{array} \right] \times 2$
12×12	ReLU

Table 4 | Final performance of agents in different settings of MuJoBan. The values reflect success ratios over the last 5000 episodes at the end of 1.5e10 training steps, averaged over agents trained with different seeds.

Setting	Agent	Level 1	Level 2	Level 3	Level 4	Level 5
Without grid pegs	Expert Planner	100.0%	100.0%	99.9%	99.0%	95.0%
Without grid pegs	Random Planner	100.0%	100.0%	99.9%	99.5%	97.7%
Without grid pegs	Vanilla Agent	100.0%	99.5%	96.2%	86.2%	62.9%
With grid pegs	Expert Planner	100.0%	99.8%	94.8%	78.2%	54.6%
With grid pegs	Random Planner	75.2%	52.0%	48.7%	42.0%	30.2%
With grid pegs	Vanilla Agent	99.8%	97.2%	68.4%	31.5%	9.4%

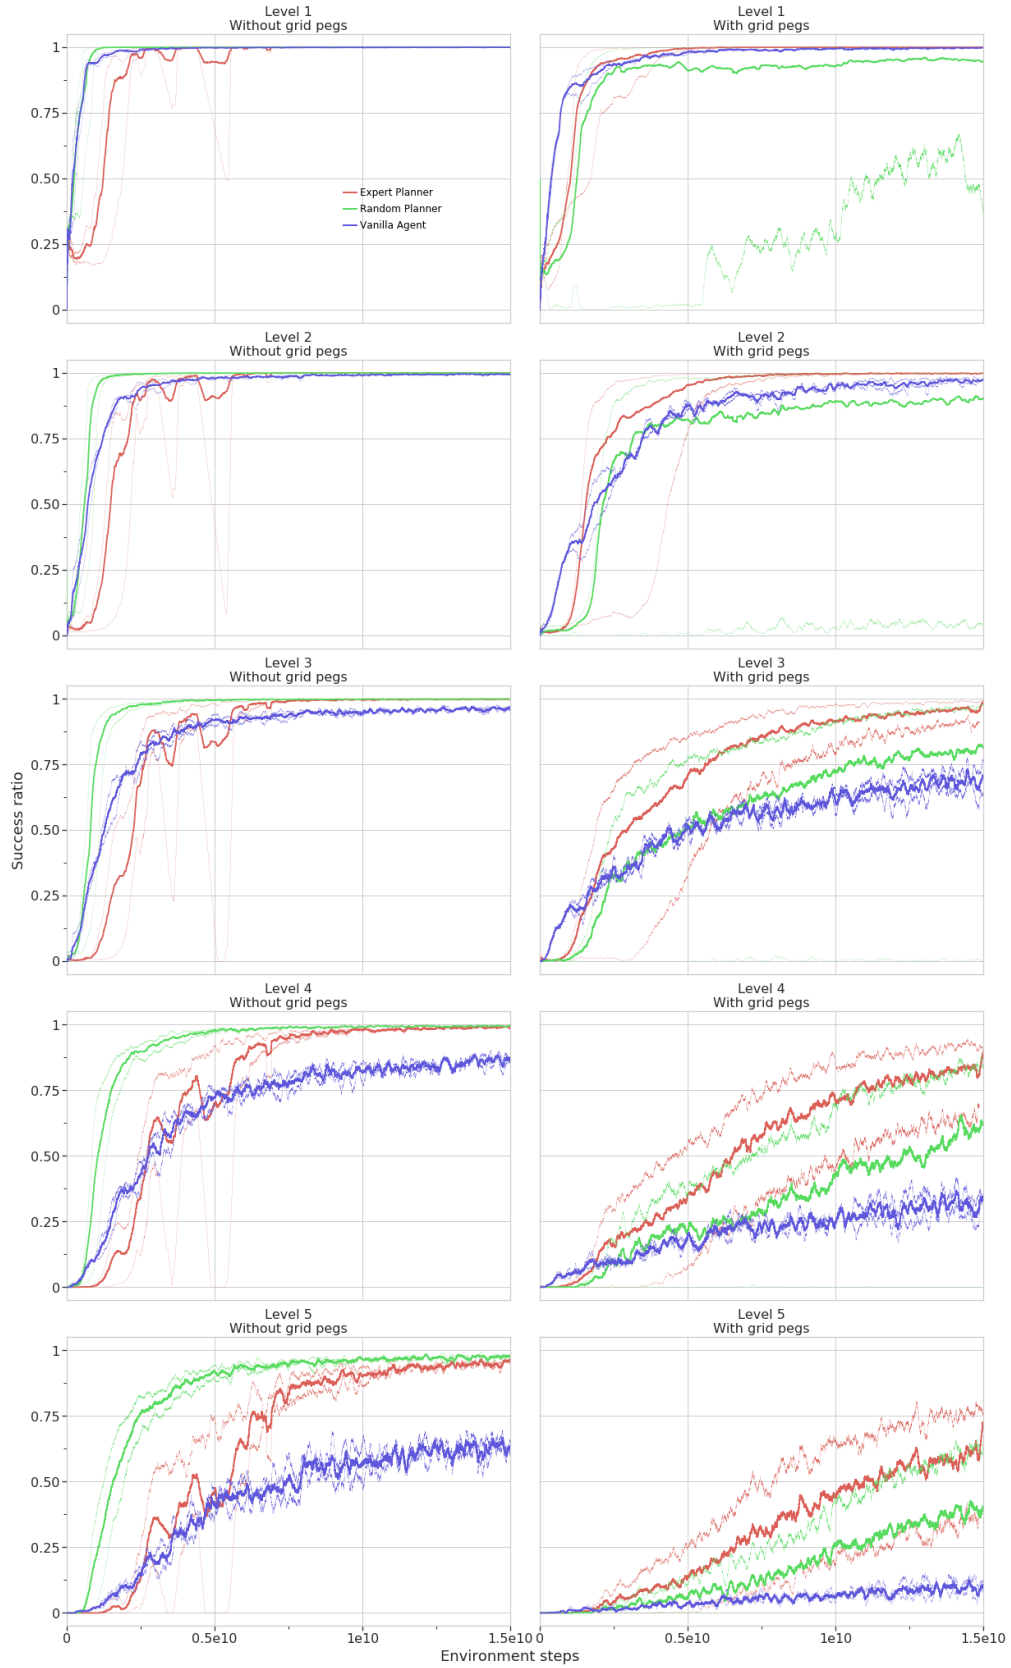


Figure 9 | Learning curves for all MuJoBan difficulty levels.