

# PDDL Gym: Gym Environments from PDDL Problems

Tom Silver and Rohan Chitnis

MIT Computer Science and Artificial Intelligence Laboratory  
tslvr@mit.edu, ronuchit@mit.edu

## Abstract

We present PDDL Gym, a framework that automatically constructs OpenAI Gym environments from PDDL domains and problems. Observations and actions in PDDL Gym are relational, making the framework particularly well-suited for research in relational reinforcement learning and relational sequential decision-making. PDDL Gym is also useful as a generic framework for rapidly building numerous, diverse benchmarks from a concise and familiar specification language. We discuss design decisions and implementation details, and also illustrate empirical variations between the 20 built-in environments in terms of planning and model-learning difficulty. We hope that PDDL Gym will facilitate bridge-building between the reinforcement learning community (from which Gym emerged) and the AI planning community (which produced PDDL). We look forward to gathering feedback from all those interested and expanding the set of available environments and features accordingly.

## 1 Introduction

The creation of benchmarks has often accelerated research progress in various subdomains of artificial intelligence (Deng et al. 2009; Wang et al. 2018; Wu et al. 2018). In sequential decision-making tasks, tremendous progress has been catalyzed by benchmarks such as the environments in OpenAI Gym (Brockman et al. 2016) and the planning tasks in the International Planning Competition (IPC) (Valati et al. 2015). Gym defines a standardized way for an agent to interact with an environment, allowing easy comparison of various reinforcement learning algorithms. IPC provides a set of planning domains and problems written in the Planning Domain Definition Language (PDDL) (McDermott et al. 1998), allowing easy comparison of symbolic planners.

In this work, we present PDDL Gym, an open-source framework that combines elements of Gym and PDDL. **PDDL Gym is a Python library that automatically creates Gym environments from PDDL domain and problem files.** The library is available at: <https://github.com/tomsilver/pddl-gym>. Pull requests are welcome!

As with Gym, PDDL Gym allows for episodic, closed-loop interaction between the agent and the environment; ICAPS Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning (PRL), 2020.




Figure 1: **Some examples of environments implemented in PDDL Gym.** From top left: Sokoban, Hanoi, Blocks, Travelling Salesman (TSP), Slide Tile, and Crafting.

the agent receives an observation from the environment and gives back an action, repeating this loop until the end of an episode. As in PDDL, PDDL Gym is fundamentally relational: observations are sets of ground relations over objects (e.g. `on(plate, table)`), and actions are templates ground with objects (e.g. `pick(plate)`). PDDL Gym is therefore particularly well-suited for relational learning and sequential decision-making research. See Figure 1 for renderings of some environments currently implemented in PDDL Gym, and Figure 2 for code examples.

The Gym API used in reinforcement learning defines a hard boundary between the agent and the environment. In particular, the agent *only* interacts with the environment by taking actions and receiving observations. The environment implements a function `step` that advances the state given an action by the agent; `step` defines the transition model of the environment. Likewise, a PDDL domain encodes a transition model via its operators. However, in typical usage, PDDL is understood to exist entirely in the “mind” of the agent. A separate process is then responsible for turning plans into actions that the agent can execute in the world.

PDDL Gym defies this convention: in PDDL Gym, PDDL domains and problems lie firmly on the environment side of the agent-environment boundary. The environment uses the PDDL files to implement the `step` function that advances the state given an action. PDDL Gym is thus perhaps best understood as a *repurposing* of PDDL. Implementation-wise, this repurposing has subtle but important implications, discussed in (§2.2) and Appendix A.

<pre>(:action put-down   :parameters (?x - block)   :precondition (and     (holding ?x)   )   :effect (and     (not (holding ?x))     (clear ?x)     (handempty)     (ontable ?x)   ) )</pre> <p style="text-align: center;">A</p>	<pre>(:init   (clear B1)   (clear B2)   (ontable B1)   (ontable B2)   (handempty) ) (:goal   (on B1 B2) )</pre> <p style="text-align: center;">B</p>	<pre>import pddlgyim  env = pddlgyim.make("PDDLEnvBlocks-v0") obs, info = env.reset() img = env.render() # We can take random actions action = env.action_space.sample(obs) obs, reward, done, info = env.step(action) # Or execute a plan plan = run_planner(...) for action in plan:   env.step(action)</pre> <p style="text-align: center;">C</p>
--	--	--

Figure 2: **PDDL Gym code examples.** A PDDL Gym environment is characterized by a PDDL domain file and a list of PDDL problem files. (A) One operator in the PDDL domain file for Blocks. (B) An excerpt of a simple PDDL problem file for Blocks. (C) After the PDDL domain and problem files have been used to register an environment with name “PDDLEnvBlocks-v0,” we can interact with this PDDL Gym environment in just a few lines of Python.

PDDL Gym serves three main purposes:

(1) *Facilitate the creation of numerous, diverse benchmarks for sequential decision-making in relational domains.* PDDL Gym allows tasks to be defined in PDDL, automatically building a Gym environment from PDDL files. PDDL offers a compact symbolic language for describing domains, which might otherwise be cumbersome and repetitive to define directly via the Gym API.

(2) *Bridge reinforcement learning and planning research.* PDDL Gym makes it easy for planning researchers and machine learning researchers to test their methods on the exact same set of benchmarks, and to develop techniques that draw on the strengths of both families of approaches. Furthermore, since PDDL Gym includes built-in domains and problems, it is straightforward to perform apples-to-apples comparisons without having to collect third-party code from disparate sources (see also (Muise 2016)).

(3) *Catalyze research on sequential decision-making in relational domains.* In our own research, we have found PDDL Gym to be very useful while studying exploration for lifted operator learning (Chitnis et al. 2020), hierarchical goal-conditioned policy learning (Silver et al. 2020a), and state abstraction (Silver et al. 2020b). Other open research problems that may benefit from using PDDL Gym include relational reinforcement learning (Lang, Toussaint, and Kersting 2012; Džeroski, De Raedt, and Driessens 2001; Tadepalli, Givan, and Driessens 2004), learning symbolic descriptions of operators (Lang, Toussaint, and Kersting 2012; Amir and Chang 2008; Pasula, Zettlemoyer, and Kaelbling 2007), discovering relational transition rules for efficient planning (Xia et al. 2019; Lang and Toussaint 2010), and learning lifted options (Konidaris, Kaelbling, and Lozano-Perez 2014; Stolle and Precup 2002; Precup, Sutton, and Singh 1998; Chentanez, Barto, and Singh 2005).

## 2 Design and Implementation

The Gym API defines environments as Python classes with three essential methods: `__init__`, which initializes the environment; `reset`, which starts a new episode and returns

an observation; and `step`, which takes an action from the agent, advances the current state, and returns an observation, reward, a Boolean indicating whether the episode is complete, and optional debugging information. The API also includes other minor methods, e.g., to handle rendering and random seeding. Finally, Gym environments are required to implement an `action_space`, which represents the space of possible actions, and an `observation_space`, which represents the space of possible observations. We next give a brief overview of PDDL files, and then we describe how action and observation spaces are defined in PDDL Gym. Subsequently, we move to a discussion of our implementation of the three essential methods. For implementation details regarding the main data structures used in PDDL Gym, see `structs.py` in the code.

### 2.1 Background: Domain and Problem Files

There are two types of PDDL files: domain files and problem files. A single *benchmark* is characterized by one domain file and multiple problem files.

A PDDL domain file includes *predicates* — named relations with placeholder variables such as  $(\text{on } ?x ?y)$  — and *operators*. An operator is composed of a name, a list of parameters, a first-order logic formula over the parameters describing the operator’s preconditions, and a first-order logic formula over the parameters describing the operator’s effects. The forms of the precondition and effect formulas are typically restricted depending on the version of PDDL. Early versions of PDDL only permit conjunctions of ground predicates (Fikes and Nilsson 1971); later versions also allow disjunctions and quantifiers (Pednault 1989). See Figure 2A for an example of a PDDL operator.

A PDDL problem file includes a set of *objects* (named entities), an *initial state*, and a *goal*. The initial state is a set of predicates ground with the objects. Any ground predicates not in the state are assumed to be false, following the closed-world assumption. The goal is a first-order logical formula over the objects (the form of the goal is limited by the PDDL version, like for operators’ preconditions and effects). Note

that PDDL (and PDDLgym) also allows objects and variables to be *typed*. See Figure 2B for a partial example.

## 2.2 Observation and Action Spaces

Each observation `obs` in PDDLgym has three components, mirroring the components of a PDDL problem file: `obs.objects` is a set containing all objects present in the problem; `obs.goal` contains the problem goal; and `obs.literals` is a set of all ground predicates that are true in the current state. These observations fully encapsulate the state of the environment, i.e., PDDLgym environments are fully-observed. The observation space is the powerset of all possible ground predicates, together with the objects and goal, which are static. This powerset is typically enormous; fortunately, it usually does not need to be explicitly computed. We expect that most algorithms for solving PDDLgym tasks will not be sensitive to its size.

The action space for a PDDLgym environment is one of the more subtle aspects of the overall framework, and there are two possible avenues to take. Instructions for taking both avenues are provided in the repository’s README, in the “Step 3: Register Gym environment” section. We include a detailed discussion of the action spaces and the related design choices in Appendix A.

## 2.3 Initializing and Resetting an Environment

A PDDLgym environment is parameterized by a PDDL domain file and a list of PDDL problem files. For research convenience, each PDDLgym environment is associated with a *test* version of that environment, where the domain file is identical but the problem files are different (for instance, they could encode more complicated planning tasks, to measure generalizability). During environment initialization, all of the PDDL files are parsed into Python objects; we use a custom PDDL parser for this purpose. When `reset` is called, a single problem instance is randomly selected.<sup>1</sup> The initial state of that problem instance is the state of the environment. For convenience, `reset` also returns (in the debugging information) paths to the PDDL domain and problem file of the current episode. This makes it easy to run to a symbolic planner and execute resulting plans in the environment; see the repository’s README for an example that uses Fast-Forward (Hoffmann 2001).

## 2.4 Implementing `step`

The `step` method of a PDDLgym environment takes in an action, updates the environment state, and returns an observation, reward, done Boolean, and debugging information. To determine the state update, PDDLgym checks whether any PDDL operator’s preconditions are satisfied given the current state and action. Note that it is impossible to “accidentally match” to an undesired operator: each operator has a unique precondition as illustrated in Figure 4C, which is generated automatically based on the passed-in action. Since actions are distinct from operators

<sup>1</sup>Problem selection when resetting an episode is the only use of randomness in PDDLgym (aside from stochastic transitions).

(Appendix A), this precondition satisfaction check is non-trivial; non-free parameters must be bound. We have implemented two inference back-ends to perform this check. The first is a Python implementation of typed SLD resolution, which is the default choice when the query involves only conjunctions. The second is a wrapper around SWI Prolog (Wielemaker et al. 2012), which permits us to handle more sophisticated preconditions involving disjunctions and quantifiers. The latter is slower, but more general, than the former. When no preconditions hold for a given action, the state remains unchanged by default. In some applications, it may be preferable to raise an error if no preconditions hold; the optional initialization parameter `raise_error_on_invalid_action` permits this.

Rewards in PDDLgym are sparse and binary. In particular, the reward is 1.0 when the problem goal is satisfied and 0.0 otherwise. Similarly, the done Boolean is True when the goal is reached and False otherwise. (In practice, a maximum episode length is often used.)

If the underlying PDDL domain has probabilistic effects, as in PPDDL (Bryce and Buffet 2008), the `step` method will parse this appropriately and choose an effect based on the given probability distribution. If the given probabilities do not sum to 1, a default trivial effect is added in.

## 2.5 Development Status

In terms of lines of code, the bulk of PDDLgym is dedicated to PDDL file parsing and inference (used in `step`). We are continuing to develop both of these features so that a wider range of PDDL domains are supported. Aspects of PDDL 1.2 that are supported by PDDLgym include STRIPS, hierarchical typing, equality, quantifiers, constants, and derived predicates. Notable features that are not supported include conditional effects and action costs. Aspects of later PDDL versions, such as numeric fluents, are not supported. Our short-term objective is to provide full support for PDDL 1.2. We have found that a wide range of standard PDDL domains are already well-supported by PDDLgym; see (§3) for an overview. We welcome requests for features and extensions, via either issues created on the Github page or email. The authors’ emails are provided at the top of this document.

# 3 PDDLgym by the Numbers

In this section, we start with an overview of the domains built into PDDLgym, as of the last date this report was updated (September 22, 2020). We then provide some experimental results that give insight into the variation between these domains, in terms of planning and model-learning difficulty. All experiments are performed on a single laptop with 32GB RAM and a 2.9GHz Intel Core i9 processor.

## 3.1 Overview of Environments

There are currently 20 domains built into PDDLgym. Most of the domains are adapted from existing PDDL repositories; the remainder are ones we found to be useful benchmarks in our own research. We have implemented custom rendering for 11 of the domains (see Figure 1 for examples). Table 1 gives a list of all environments, their sources, and

Domain Name	Source	Rendering Included	Probabilistic	Average FPS
Baking	Ours	No	No	5897
Blocks	(Helmert 2011)	Yes	No	7064
Casino	Ours	No	No	7747
Crafting	Ours	Yes	No	4568
Depot	(Helmert 2011)	No	No	97
Doors	(Konidaris and Barto 2007)	Yes	No	917
Elevator	(Helmert 2011)	No	No	3501
Exploding Blocks	(Bryce and Buffet 2008)	Yes	Yes	6260
Ferry	(CSU 2002)	No	No	1679
Gripper	(Soar 2020)	Yes	No	319
Hanoi	(Soar 2020)	Yes	No	4580
Meet-Pass	(CSU 2002)	No	No	7380
Rearrangement	Ours	Yes	No	3808
River	(Bryce and Buffet 2008)	No	Yes	18632
Search and Rescue	Ours	Yes	No	3223
Slide Tile	(Soar 2020)	Yes	No	3401
Sokoban	(Helmert 2011)	Yes	No	155
Triangle Tireworld	(Bryce and Buffet 2008)	No	Yes	6491
TSP	(Soar 2020)	Yes	No	1688
USA Travel	Ours	No	No	1251

Table 1: List of the 20 domains currently included in PDDLgym, as of the last date this report was updated (September 22, 2020). For each environment, we report the original source of the PDDL files, whether or not we have implemented custom rendering, whether or not the domain has probabilistic effects, and the average frames per second (FPS). The FPS is calculated by executing a random policy for 100 episodes of 10 timesteps each, with no rendering.

their average frames per second (FPS) calculated by executing a random policy for 100 episodes of 10 timesteps each.

### 3.2 Variation in Environment Difficulty

We now provide some results illustrating the variation between the domains built into PDDLgym. We examine two axes of variation: planning difficulty and difficulty of learning the transition model.

Figure 3 (left) illustrates the average time taken by Fast-Forward (Hoffmann 2001) to find a plan in each of the deterministic environments, averaged across all problem instances. The results reveal a considerable range in planning time, with the most difficult domain (Depot, omitted from the plot for visual clarity) requiring two orders of magnitude more time than the simplest one (TSP). The results also indicate that many included domains are relatively “easy” from a modern planning perspective. However, even in these simple domains, there are many interesting challenges to be tackled, such as learning the true PDDL operators from interaction data, or defining good state abstractions amenable to learning. One can always make larger problem instances if desired, to push the limits of modern planners.

Figure 3 (right) provides insight into the difficulty of learning transition models in some of the environments. For each environment, an agent executes a random policy for episodes of horizon 25. The observed transitions are used to learn transition models, which are then used for planning on a suite of test problems. The fraction of test problems solved is reported as an indicator of the learned transition model. To learn the transition models, we use first-order logic decision tree (FOLDT) learning (Blockeel and De Raedt 1998).

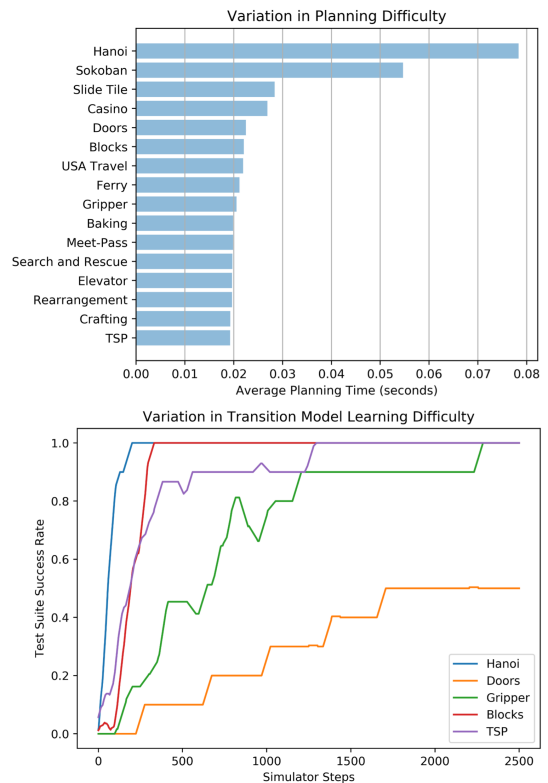


Figure 3: **Variation among PDDLgym environments.** The PDDL domains and problems built into PDDLgym vary considerably in terms of planning difficulty (top) and model learning difficulty (bottom). See text for details.

Five domains are visualized for clarity; among the remaining ones, several are comparable to the ones shown, but others, including Baking, Depot, and Sokoban, are difficult for our learning method: FOLDT learning is unable to find a model that fits the data in a reasonable amount of time. Of course, model-learning difficulty varies considerably with the learning method and the exploration strategy. We have implemented simple strategies here to show these results, but these avenues for future research are exactly the kind that we hope to enable with PDDLgym.

## 4 Conclusion and Future Work

We have presented PDDLgym, an open-source Python framework that automatically creates OpenAI Gym environments from PDDL domain and problem files. Our empirical results demonstrate considerable diversity among the built-in environments. We have been using PDDLgym actively in our own research on relational sequential decision-making and reinforcement learning. We also hope to interface PDDLgym with other related open-source frameworks, particularly the PDDL collection in planning.domains (Muise 2016), so that one can use PDDLgym simply by specifying a URL (along with information about free parameters).

We look forward to gathering feedback from the community and expanding the set of environments and features.

## References

- Amir, E., and Chang, A. 2008. Learning partially observable deterministic action models. *Journal of Artificial Intelligence Research* 33:349–402.
- Blockeel, H., and De Raedt, L. 1998. Top-down induction of first-order logical decision trees. *Artificial intelligence* 101(1-2):285–297.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI gym. *arXiv preprint arXiv:1606.01540*.
- Bryce, D., and Buffet, O. 2008. International planning competition uncertainty part: Benchmarks and results. In *In Proceedings of IPC*. Citeseer.
- Chentanez, N.; Barto, A. G.; and Singh, S. P. 2005. Intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, 1281–1288.
- Chitnis, R.; Silver, T.; Tenenbaum, J.; Kaelbling, L. P.; and Lozano-Perez, T. 2020. GLIB: Exploration via goal-literal babbling for lifted operator learning. *arXiv preprint arXiv:2001.08299*.
- CSU. 2002. Colorado state university PDDL repository. <https://www.cs.colostate.edu/meps/repository.html>.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Džeroski, S.; De Raedt, L.; and Driessens, K. 2001. Relational reinforcement learning. *Machine learning* 43(1-2):7–52.
- Fikes, R. E., and Nilsson, N. J. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial intelligence* 2(3-4):189–208.
- Helmert, M. 2006. The fast downward planning system. *Journal of Artificial Intelligence Research* 26:191–246.
- Helmert, M. 2011. Pyperplan. <https://bitbucket.org/malte/pyperplan>.
- Hoffmann, J. 2001. FF: The fast-forward planning system. *AI magazine* 22(3):57–57.
- Konidaris, G., and Barto, A. G. 2007. Building portable options: Skill transfer in reinforcement learning. In *IJCAI*, volume 7, 895–900.
- Konidaris, G.; Kaelbling, L.; and Lozano-Perez, T. 2014. Constructing symbolic representations for high-level planning. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.
- Lang, T., and Toussaint, M. 2010. Planning with noisy probabilistic relational rules. *Journal of Artificial Intelligence Research* 39:1–49.
- Lang, T.; Toussaint, M.; and Kersting, K. 2012. Exploration in relational domains for model-based reinforcement learning. *Journal of Machine Learning Research* 13(Dec):3725–3768.
- McDermott, D.; Ghallab, M.; Howe, A.; Knoblock, C.; Ram, A.; Veloso, M.; Weld, D.; and Wilkins, D. 1998. Pddl-the planning domain definition language.
- Muise, C. 2016. Planning.Domains. In *The 26th International Conference on Automated Planning and Scheduling - Demonstrations*.
- Pasula, H. M.; Zettlemoyer, L. S.; and Kaelbling, L. P. 2007. Learning symbolic models of stochastic domains. *Journal of Artificial Intelligence Research* 29:309–352.
- Pednault, E. P. 1989. ADL: Exploring the middle ground between strips and the situation calculus. *Kr* 89:324–332.
- Precup, D.; Sutton, R. S.; and Singh, S. 1998. Theoretical results on reinforcement learning with temporally abstract options. In *European conference on machine learning*, 382–393. Springer.
- Silver, T.; Chitnis, R.; Ajay, A.; Tenenbaum, J.; and Kaelbling, L. P. 2020a. Learning skill hierarchies from predicate descriptions and self-supervision. In *AAAI GenPlan Workshop*.
- Silver, T.; Chitnis, R.; Curtis, A.; Tenenbaum, J.; Lozano-Perez, T.; and Kaelbling, L. P. 2020b. Planning with learned object importance in large problem instances using graph neural networks. *arXiv preprint arXiv:2009.05613*.
- Soar. 2020. Soar group PDDL repository. <https://github.com/SoarGroup>.
- Stolle, M., and Precup, D. 2002. Learning options in reinforcement learning. In *International Symposium on abstraction, reformulation, and approximation*, 212–223. Springer.
- Tadepalli, P.; Givan, R.; and Driessens, K. 2004. Relational reinforcement learning: An overview. In *Proceedings of the ICML-2004 workshop on relational reinforcement learning*, 1–9.
- Vallati, M.; Chrapa, L.; Grześ, M.; McCluskey, T. L.; Roberts, M.; Sanner, S.; et al. 2015. The 2014 international planning competition: Progress and trends. *Ai Magazine* 36(3):90–98.
- Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Wielemaker, J.; Schrijvers, T.; Triska, M.; and Lager, T. 2012. SWI-Prolog. *Theory and Practice of Logic Programming* 12(1-2):67–96.
- Wu, Z.; Ramsundar, B.; Feinberg, E. N.; Gomes, J.; Geniesse, C.; Pappu, A. S.; Leswing, K.; and Pande, V. 2018. MoleculeNet: a benchmark for molecular machine learning. *Chemical science* 9(2):513–530.
- Xia, V.; Wang, Z.; Allen, K.; Silver, T.; and Kaelbling, L. P. 2019. Learning sparse relational transition models. *International Conference on Learning Representations*.



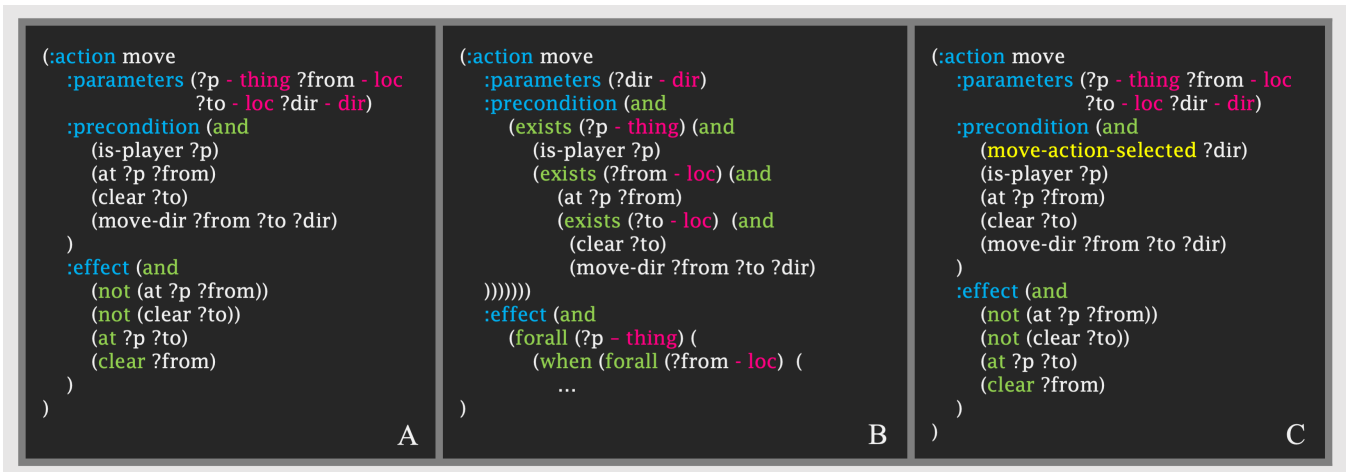


Figure 4: **Explicating free parameters in PDDL operators.** PDDL operators traditionally conflate free and non-free parameters. For example, in a typical move operator for Sokoban (A), the free parameter `?dir` is included alongside non-free parameters. PDDLGym must distinguish free parameters to properly define the action space. One option would be to require that all operator parameters are free, and introduce quantifiers in the operator body accordingly (B); however, this is cumbersome and leads to clunky, deeply nested operators, so we do *not* do this. Instead, we opt to introduce new predicates that are tied to operators, and whose parameters are just the operators’ free parameters (C). An example of such a new predicate is shown in yellow (`move-action-selected`).

## A Action Space Details

The action space for a PDDLGym environment is one of the more subtle aspects of the overall framework, and there are two possible avenues to take. Instructions for taking both avenues are provided in the repository’s README, in the “Step 3: Register Gym environment” section.

The first avenue is appropriate if one wants to simply use off-the-shelf PDDL files with PDDLGym. One can do so by setting `operators_as_actions` to `True` in the environment registration, which tells PDDLGym that the operators present in the PDDL domain file should themselves be treated as the actions in the environment, parameterized by those operators’ parameters.

The second avenue is recommended for more serious research, and stems from the semantic difference between “operators” in classical AI planning and “actions” in reinforcement learning. In AI planning, actions are typically equated with ground operators — operators whose parameters are bound to objects. However, in most PDDL domains, only some operator parameters are *free* (in terms of controlling the agent); the remaining parameters are included in the operator because they are part of the precondition/effect expressions, but can be derived from the current state or the choice of free parameters. PDDL itself makes no distinction between free and non-free parameters. For example, consider the operator for Sokoban shown in Figure 4A. This operator represents the rules for a player (`?p`) moving in some direction (`?dir`) from one cell (`?from`) to another cell (`?to`). In a real game of Sokoban, the only choice that an agent makes is what direction to move — only the `?dir` parameter is free. The player `?p` is always the same, `?from` is defined by the agent’s location in the current state, and

`?to` can be derived from `?from` and the agent’s choice of `?dir`. To properly define the action space for a PDDLGym environment, we must explicitly distinguish free parameters from non-free ones. One option is to require that operator parameters are all free. Non-free parameters could then be folded into the preconditions and effects using quantifiers (Pednault 1989); see Figure 4B for an example. However, this is cumbersome and leads to clunky, deeply nested operators. Instead, we opt to introduce new predicates that represent operators, and whose variables are these operators’ free parameters. We then include these predicates in the preconditions of the respective operators; see Figure 4C for an example. Doing so requires only minimal changes to existing PDDL files and does not affect readability, but requires adding in domain knowledge about the agent-environment boundary. Note that this domain knowledge is equivalent to defining an action space, which is very commonly done in reinforcement learning and is not a strong assumption. In this case, the action space of a PDDLGym environment is a discrete space over all possible groundings of the newly introduced predicates.

When sampling from the action space of a PDDLGym environment, PDDLGym will automatically only sample *valid* actions, i.e., actions that satisfy the preconditions of some operator. This check for validity is done using Fast Downward’s translator (Helmert 2006), which can add non-negligible overhead in large problem instances.