

# Planning Domain Simulation: An Interactive System for Plan Visualisation

Emanuele De Pellegrin, Ronald P. A. Petrick

Edinburgh Centre for Robotics  
 Heriot-Watt University  
 Edinburgh, Scotland, United Kingdom  
 ed50@hw.ac.uk, R.Petrick@hw.ac.uk

## Abstract

Representing and manipulating domain knowledge is essential for developing systems that can visualize plans. This paper presents a novel plan visualisation system called Planning Domain Simulation (PDSim) that employs knowledge representation and manipulation techniques to support the plan visualization process. PDSim can use PDDL or the Unified Planning Library’s Python representation as the underlying language for modelling planning problems and provides an interface for users to manipulate this representation through interaction with the Unity game engine and a set of planners. The system’s features include visualising plan components, and their relationships, identifying plan conflicts, and examples applied to real-world problems. The benefits and limitations of PDSim are also discussed, highlighting future research directions in the area.

## Introduction

Modelling planning domains that are both correct and robust can be a challenging problem, especially in real-world domains. For instance, consider the following robot planning task: a set of robots are deployed in a factory to help with warehouse logistics. The robots can navigate on a pre-defined grid map with simple 4-way movements, pick up and drop boxes, and deliver objects to a van parked in the warehouse. The problem also imposes certain limitations: the robots cannot cross each other and the vans can only accept a specific box. The above problem could be viewed as a slightly modified version of the sequential *Floor Tile* domain from the 2011 International Planning Competition (IPC),<sup>1</sup> a decision-making problem inspired by a real-world scenario that can be modelled using a representation language such as PDDL (McDermott et al. 1998). From a representation point of view, the grid could be modelled as a set of interconnected nodes denoting locations in the warehouse for objects and agents (e.g., vans, boxes, and robots), as illustrated in Figure 1. A simple goal might be to ensure that particular objects are in specific locations (e.g., *box1* is in *van1*).

Using the above model, we can quickly find a *valid* solution to the problem using classical automated planning tech-

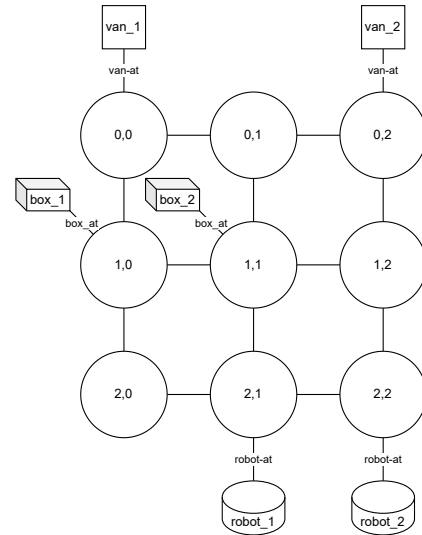


Figure 1: Initial state for the Warehouse problem.

LEFT (R1, C2-1, C2-0)	LEFT (R1, C2-1, C2-0)
UP (R1, C2-0, C1-0)	UP (R1, C2-0, C1-0)
PICKUP (R1, B1, C1-0)	UP (R1, C1-0, C0-0)
UP (R1, C1-0, C0-0)	PICKUP (R1, B1, C1-0)
LOAD (R1, B1, C0-0, V1)	LOAD (R1, B1, C0-0, V1)

Figure 2: Example plan outputs for the Warehouse problem. The plan on the left is correct, while the plan on the right is found after introducing an error in the domain.

niques. For instance, Figure 2 (left) shows a plan generated by the FastDownward planner (Helmert 2006) for the problem in Figure 1, where a robot moves to grid `cell_1_0` to pick up the box before delivering it to the van at `cell_0_0`.

Figure 2 (right) shows an alternative action sequence, generated using an incorrect version of the domain. Although the plan is similar to the one on the left, it is *incorrect*: the robot executes the pickup action when in grid `cell_0_0` before loading the van. (This plan is the result of a missing precondition on the pickup action which normally ensures that the robot and object are in the same cell). While this kind of error can be trivial to debug and correct by an expert

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup><https://github.com/potassco/pddl-instances/tree/master/ipc-2011/domains/floor-tile-sequential-satisficing>

knowledge engineer, this isn't always the case for novices using languages such as PDDL. Catching modelling errors, such as incorrect logic in action preconditions and effects or missing properties in the initial state, can still be difficult due to the complexity of the knowledge that needs to be specified and the level of abstraction that is often required for ensuring the generation of tractable solutions.

In this paper, we present the Planning Domain Simulation (PDSim) system (De Pellegrin 2020; De Pellegrin and Petrick 2021, 2022, 2023), a framework for visualising and simulating a range of planning problems (classical, numerical, or temporal) using the Unified Planning Library (UPL) (Micheli and Bit-Monnot 2022) and the Unity game engine (Unity Technologies 2022). Using UPL or PDDL, the user can define the domain knowledge and the problem formulation (e.g., planner requirements, types and objects, plus standard definitions of the domain and problem). A planner then uses this information to check that a solution exists and to generate a plan that satisfies the goal. Using the generated plan, PDSim interprets the action effects as 3D animations and graphics effects in Unity to deliver a visual representation of the world and its actions during plan execution, which can aid the user in assessing the validity of the plan.

While several tools already exist to aid in the process of validating planning models—notably plan validation tools like VAL (Howey and Long 2003) and formal plan verification methods such as (Bensalem, Havelund, and Orlandini 2014; Cimatti, Micheli, and Roveri 2017; Hill, Komentanskaya, and Petrick 2020)—approaches based on visual simulation and visual feedback can also play an important role in addressing this problem: visual tools can serve as powerful environments for displaying, inspecting, and simulating the planning process, which can aid in plan explainability for human users (Fox, Long, and Magazzeni 2017).

In this paper, we describe the structure, components, and features of PDSim that are responsible for providing visualisations, and illustrate how PDSim can be used to simulate planning problems. PDSim is built by extending the Unity game engine editor, and can use the components offered by the engine such as a path planner, scene management, and visual scripting, among others. The system uses a backend server that is responsible for defining planning problems using UPL or PDDL, managing plan generation and problem compilation, and providing support for a range of modelling features including typing, temporal actions, and action costs.

The rest of the paper is organised as follows. First, we give an overview of automated planning with PDDL and review work related to plan visualisation. We then describe how knowledge is represented in PDSim and outline the structure of PDSim's main components. Examples are provided for a number of planning domains. Finally, we conclude with future work and planned extensions to PDSim.

## Background and Related Work

### Automated Planning with PDDL

Automated planning is a decision-making task that involves reasoning about the sequence of actions (a plan) that achieves a set of goals (Ghallab, Nau, and Traverso 2004;

Haslum et al. 2019). A planning problem  $\Pi$  can be thought of as a tuple  $\Pi = \langle P, A, I, G \rangle$ , where  $P$  is a set of properties that define a state space (including possibly a set of objects),  $A$  is a set of actions,  $I$  is a set of initial state properties, and  $G$  is the set of goal conditions to be achieved. It is useful to think of a planning problem as a state transition system, where a state captures all the properties that are true at some point in time, and actions transition states to new states. A solution to the planning problem is a sequence of actions, called a plan, that when applied to the initial state  $I$  transitions to a state in which the goal conditions  $G$  are true.

Automated planning has been used in a variety of applications such as robotics, video games, logistics, and natural language processing. Intuitively, planning can be thought of as a search process that enables an autonomous agent to generate a plan to achieve its goals. In this view, plan generation typically involves the following steps:

1. **Problem Definition:** Specifying the planning model  $\Pi$  (properties, actions, initial state, and goals) that captures the operating environment of the agent.
2. **Search Space Generation:** Creating a representation of the possible states that can be achieved by applying actions from the initial state to the goal state.
3. **Search:** Applying a search algorithm that explores the state space and selects a plan that satisfies the goal.

Planning problems are composed of two parts: the domain definition which specifies the state properties and actions, and the problem definition which specifies the initial state and the goal. State properties are specified using (parameterized) predicates that can be true or false in a given state and that capture attributes of the environment, objects, or agents. For instance, `(clear cell_0_1)` might denote that location (0,1) is empty, `(at box1 cell_1_1)` might capture the fact that box1 is at location (1,1), and `(robot-empty robot1)` might represent the idea that robot1 isn't carrying anything. Predicates specify the initial state and goal conditions of the planning problem, and are also used to describe the preconditions and effects of actions.

Actions are formalised using a schema that specifies the parameters, preconditions, and effects of each action, as in Figure 3 using PDDL or Figure 4 using UPL and Python. The preconditions capture the conditions that must be true in a state to perform the action, while the effects describe the state changes after an action is performed. For instance, the `load-truck` action in Figures 3 and 4 has three parameters: a package ( $?p$ ), a truck ( $?t$ ), and a location ( $?l$ ). A package  $?p$  can be loaded onto a truck  $?t$  provided  $?t$  is at location  $?l$ , `(at ?t ?l)`, and  $?p$  is at location  $?l$ , `(at ?p ?l)`. As a result of applying `load-truck`, the package will no longer be at  $?l$ , `(not (at ?p ?l))`, and will be in the truck, `(in ?p ?t)`. When an action is chosen by the planner to be part of the plan, its parameters will be replaced by objects in the planning problem (e.g., `van1` for  $?t$ , `box2` for  $?p$ , and `cell_1_2` for  $?l$ ).

Domain and problem definitions are used as input to an automated planner that can reason about the changes in the world state when actions are applied, and generate a plan that achieves the goal conditions. A plan is typically a se-

```

(:action load-truck
 :parameters (?p, ?t, ?l)
 :precondition (and (at ?t ?l)
                    (at ?p ?l))
 :effect (and (not (at ?p ?l))
              (in ?p ?t))
)

```

Figure 3: PDDL representation of the `load-truck` action in the Warehouse domain.

```

lt = InstantaneousAction("loadTruck",
    p=parcel, t=truck, l=location)

lt.add_precondition(at(lt.t, lt.l)
    &
    at(lt.p, lt.l))
lt.add_effect(at(lt.p, lt.l), False)
lt.add_effect(in(lt.p, lt.t), True)

```

Figure 4: Representation of the `load-truck` action in the UPL Python format.

quence of actions, as shown in Figure 2, where each row represents an action and its (grounded) parameters, and the parameters in the action schema have been replaced with objects from the problem definition.

## Plan Visualisation

PDSim (De Pellegrin and Petrick 2023) is part of the small ecosystem of simulators for automated planning which use visual cues and animations to translate the output of a plan into a 3D or 2D environment. The closest approach to ours is Planimation (Chen et al. 2020) which uses Unity as the front-end engine to display objects and animate their position in a given plan. Planimation defines animations using an ad hoc language (namely, an animation profile) similar to PDDL. This differs from PDSim, where animations are defined using Unity’s visual scripting system.<sup>2</sup>

The Logic Planning Simulator (LPS) (Tapia, San Segundo, and Artieda 2015) also provides a planning simulation system that represents PDDL objects with 3D models in a user-customisable environment. The approach is integrated with a SAT-based planner and a user interface that enables plan execution to be simulated while visualising updates to the world state and individual PDDL properties in the 3D environment. LPS is not based on Unity but provides the user with a simple interface for plan visualisation. Several user-specified files are also required to define 3D object meshes, the relationship between PDDL elements and 3D objects, and the specific animation effects.

vPlanSim (Roberts et al. 2021) is a similar application that also aims to provide a 3D visualisation of a plan but with a number of important differences. While vPlanSim offers a simple and fast custom graphical environment for creating plan simulations with few dependencies, PDSim uses

the Unity game engine to offer the user industry-standard tools for creating realistic scenarios. PDSim also provides a language-agnostic tool to set up simulations which is key for users who are not familiar with PDDL and Unity.

Table 1 highlights the main differences between PDSim, Planimation, and vPlanSim in how they tackle the problem of plan visualisation. Unlike other visualisation systems that offer a limited selection of planners, PDSim, thanks to its integration with UPL, supports a range of planners such as FastDownward, Tamer, and LPG, providing capabilities for both numerical and temporal planning. This distinguishes PDSim from existing visualization systems. Additionally, PDSim also offers compatibility with ROS for robotics applications, a feature that to the best of our knowledge is not supported by other visualization systems. In contrast, Planimation relies on an animation profile in a definition language similar to PDDL, while vPlanSim requires scripting using a Python routine. For users seeking more in-depth modifications to the visualisation, PDSim provides a C# API.

Several systems also exist to help users formalise planning domains and problems through user-friendly interfaces. For instance, GIPO (Simpson, Kitchin, and McCluskey 2007), ItSimple (Vaquero et al. 2007), and VIZ (Vodrázka and Chrapa 2010) use graphical illustrations of the domain and problem elements, removing the requirement of PDDL language knowledge to help new users approach planning domain modelling for the first time. Tools such as Web Planner (Magnaguagno et al. 2017) and Planning.Domains (Muisse 2016) use Gantt charts or tree-like visualisations to illustrate generated plans and the state spaces searched by a particular planning algorithm. Plan-Curves (Le Bras et al. 2020) uses a novel interface based on time curves (Bach et al. 2015) to display timeline-based multiagent temporal plans distorted to illustrate the similarity between states. All of these tools attempt to assist users in understanding how a plan is generated and to help detect potential errors in the modelling process.

Simulators are also prevalent in robotics applications and multiple systems make use of game engines to provide virtual environments, such as MORSE (Echeverria et al. 2011) or Drone Sim Lab (Ganoni and Mukundan 2017). Game engines also offer several benefits such as multiple rendering cameras, physics engines, realistic post-processing effects, and audio engines, without the need to implement these features from scratch (Ganoni and Mukundan 2017), making them desirable tools for simulation. For example, Unity has been used as a tool for data visualisation, architectural prototypes, robotics simulation (Green et al. 2020), synthetic data generation for computer vision (Fort, Hogins, and Davis 2021), and machine learning applications (Haas 2014; Craighead, Burke, and Murphy 2008). There are also interesting use cases of Unity related to AI and planning, including the Unity AI Planner,<sup>3</sup> an integrated planner being created by Unity as a component for developing AI solutions for videogames, and Unity’s machine learning agents,<sup>4</sup> a so-

<sup>2</sup><https://docs.unity3d.com/Packages/com.unity.visualscripting@1.7/manual/vs-nodes-reference.html>

<sup>3</sup><https://docs.unity3d.com/Packages/com.Unity.ai.planner@0.0/manual/index.html>

<sup>4</sup><https://github.com/Unity-Technologies/ml-agents>

	PDSim	Planimation	vPlanSim
<b>Supported Planners</b>	Supports a wide range of planners, (see Back-End section)	planning.domains solver API	External PDDL planner
<b>Planning Support</b>	Classical, Temporal, Numeric	Classical	Classical
<b>Visualisation</b>	Unity Editor or executable targeting user OS	Web interface	VTK and PyQt
<b>Animation Creation</b>	Visual scripting in-engine system	Animation profile definition (PDDL-like)	Scripting routines with Python
<b>Representations</b>	PDDL, ANML, Python, C#	PDDL	PDDL
<b>Miscellaneous</b>	Middleware support for robotics with ROS	Can host local versions of Planimation	Can generate PDDL problems interactively

Table 1: Comparison of plan visualisation systems.




Figure 5: Mapping between the current planning problem state description (Facts), the UPL Internal Representation used to schedule the selected planner, and the PDSim Representation using the C# counterpart for Unity.

lution for training and displaying agents whose behaviour is driven by an external machine learning component.

## Knowledge Representation

PDSim visualises states and plans in the Unity game environment. Knowledge is visualised by defining a planning problem directly using UPL or by using UPL to parse a PDDL representation. After a plan is generated, both the plan and the problem definition are converted to a protocol buffer representation<sup>5</sup> that will be later mapped to Unity’s game engine objects. In Unity, the user defines the procedures, animations, and final visualisation of the plan. In this section, we discuss the underlying planning model and how it is mapped to concepts in Unity.

## Mapping Planning Components into Unity

Unity does not have built-in support for planning problem modelling languages but instead uses C# as a scripting language. As a result, components must be mapped into C# constructs (classes) to be represented in Unity. For a given planning domain and problem description, a set of basic constructs must be translated for plan visualisation: predicates, actions, types, and constants. Figure 5 shows the mapping between different types of knowledge and how that knowledge is manipulated for visualisation in PDSim. ‘Facts’ correspond to the high-level knowledge that the user wants to represent (e.g., the configuration of the Warehouse environment). This high-level representation can be mapped to the ‘UPL Internal Representation’ by using PDDL or the Python

<sup>5</sup><https://protobuf.dev/>

library. Here knowledge is further manipulated and used to perform the search for plans using the planning components. The last block corresponds to the ‘PDSim Representation’ that maps the planning model to Unity C# components using the protocol buffer representation.

Predicates define the properties that objects have (or don’t have) in a particular state. In PDSim, predicates are encoded as Object Oriented Programming (OOP) classes. In particular, PDSim differentiates between Boolean, Numeric, or Symbolic predicates. Boolean predicates can be either true or false. The animation of Boolean predicates can be split in two ways and the user can customise the behaviour of both values. Numeric predicates represent a predicate that can hold a numeric value. Animations of Numeric predicates can represent increases or decreases in numeric assignments. Finally, Symbolic predicates map animations to predicates that have a symbolic value, such as a constant.

Actions are defined by their preconditions and effects. Actions in PDSim are also represented by classes that store the set of effects and the possible objects that can be used with the action. Types are used to define a specific property for an object, in a parent-child relationship. In C#, types are represented with a tree-type structure so that if an object is of a particular type it inherits all the possible actions that the supertype has access to. For example, a *robot* could be a *physi-Object* child type that inherits all the animation available to this type. Types are not necessary for PDSim, as a predicate animation can also be used to define the type of constant on the Unity side. For example, *cube(?c)* can be mapped to an animation that can spawn a cube model or sprite and set its position in the 3D environment. Constants are used to refer to specific objects in the planning problem. In C#, and more particularly in Unity, constants represent the virtual actors in the scene. These can be 3D or 2D models and the animations that are directly applied to them.

## Converting the Planning Model to PDSim

The planning model is converted to a protocol buffer representation that maps to an internal C# PDSim model representing the components described above (predicates, actions, constants, etc.). This model is used in Unity to set up the simulation, where domain entities such as actions, types, and predicates are used to set up the core Unity simulation.

```

(:init
  (at box1 cell_1_1)
  (at box2 cell_1_2)
  (at robot1 cell_0_0)
  (at robot2 cell_2_2)
  (robot-empty robot1)
  (robot-empty robot2)
  ...
  (up cell_0_1 cell_1_1)
  (down cell_1_1 cell_0_1)
  (right cell_0_2 cell_0_1)
  (left cell_0_1 cell_0_2)
  ...)

```

Figure 6: Example of an initial state in PDDL format.




Figure 7: Initial state representation translation in PDSim where the PDDL predicates are visualised as objects and physical positions in a 3D space.

Similarly, problem components such as constants and the initial state are used to set up a Unity-level scene. Once these components are defined, users can customise them using the Unity editor, for instance configuring multiple problems for the same domain, or multiple simulations for different plans.

Figure 6 shows the PDDL problem definition for the initial state described in the introduction. The `at` predicate is used to describe the position of a physical object (robots, boxes and vans), `robot-empty` represents if a robot is carrying a box or not, `van-request` represents which box is requested by a particular van, and `up`, `down`, `right`, `left` represent the connections between cells in the grid. The same PDDL representation can be visualized with 3D models in PDSim as shown in Figure 7. The PDDL `:init` block from Figure 6 can be animated in PDSim by assigning translation sequences to the physical objects and displaying them in game mode.

### From C# to Animations

The planning domain description is used to build the core elements and animations for the simulation. The types and objects define the visual aspect of the Unity simulation: 3D models or 2D sprites. Once mapped, predicates are used to define the 2D/3D animations using Unity’s visual scripting option. This visual scripting language is used to define and interface with Unity’s common transformation operations, path planning, audio emission, and particle effects systems.

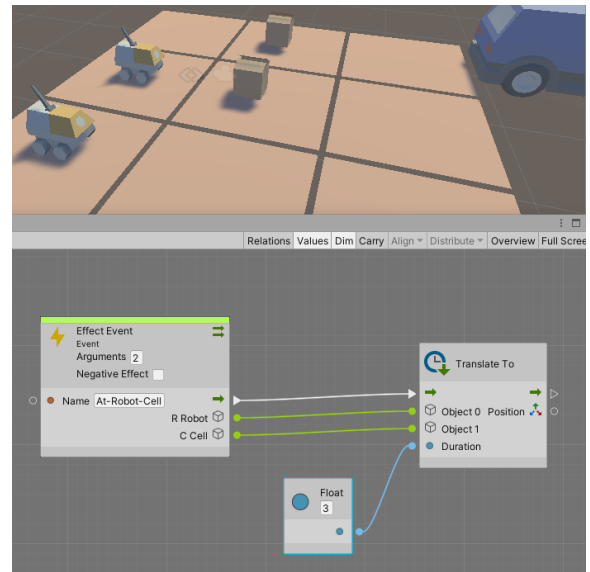


Figure 8: Example animation definition using Unity’s visual scripting graph showing the event for the ‘`at(robot,cell)`’ predicate and the node to execute (“Translate To”) when the predicate is executed in one of the plan action effects.

For example, Figure 8 shows an animation definition for the earlier Warehouse planning problem, for a predicate that captures the movement of the robot position from the current grid to an adjacent cell. Action effects are the animated components, where every predicate in the effects list that has an associated animation graph will execute an animation at simulation time.

Users can define their own behaviours in the virtual scene for every predicate they want to animate. The example in Figure 8 shows a simple translation animation from an object position to a target position. In particular, the example shows one of the custom animation nodes developed in PDSim to help simplify the creation of animations for new users. Every predicate in an action’s effect can have one of these graphs linked to it, and every graph comes with an *EffectEvent* that is invoked during plan simulation with the corresponding objects from the Unity scene (i.e., the objects in the plan’s action).

To simplify the development of new animations, and to help new users with visual scripting, a set of predefined animation nodes has been created which cover a number of useful simulation cases that frequently arise, such as:

1. **TranslateToPoint:** Move a particular object in the scene to a specific point in the world or to another object’s position (using path planning or simple interpolation).
2. **TranslateToObject:** Move a particular object in the scene to a specific other object in the world (using path planning or simple interpolation).
3. **SpawnObject:** Instantiate an object (i.e., a 3D mesh) in the scene.
4. **PlayPauseParticle:** Create and either play or pause a particle effect.

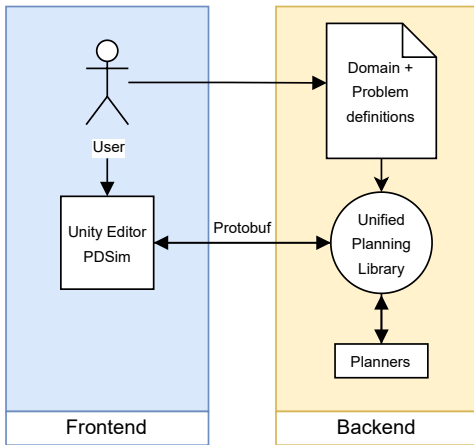


Figure 9: High-level PDSim system architecture.

5. **PlayPauseSound:** Create and either play or pause an audio effect.
6. **GetCurrentPlanActions:** Can get the current simulated actions. Multiple actions can return if the simulated plan is temporal and multiple actions are currently being executed. It also returns metadata for the action such as parameters (objects) and action duration.

### System Architecture

The high-level structure of the PDSim system is shown in Figure 9. PDSim can be imported into Unity3D as a common asset, where the Unity editor interface is used to interact with PDSim components, such as setting the simulation scene, creating animations, or importing 3D or 2D models. PDSim also relies on a Python backend implementation, which is used to parse PDDL files and generate plans. A PDSim simulation is initialised and handled by the backend server running the UPL<sup>6</sup> library, which is responsible for parsing and building a protocol buffer representation of the planning model and running a user-defined planner (defaulting to FastDownward) to generate a plan. UPL is a planner-agnostic framework for Python, which increases PDSim’s modularity and lets users select their preferred planner implementation, separating it from the simulation stage itself which comes later in the process. We describe the major components of PDSim below.

#### Front-End

Unity (Unity Technologies 2022) is a popular state-of-the-art game engine used for building 3D projects across a range of diverse applications. In PDSim, Unity provides the front-end interface and is responsible for handling all of the 2D/3D graphics and animations related to the simulation.

One of the fundamental design concepts used by Unity is the idea of *composition*, which means that an object can be *composed* of different types of objects. In particular, Unity’s component system provides the capability for every object in a Unity scene to be assigned custom scripts or modules, such as a rigid body for the physics simulation, a collision

<sup>6</sup><https://github.com/aiplan4eu/unified-planning>




Figure 10: Simulation object example with a Robot type, highlighting the customisations the user can work with, such as adding custom models and defining points in the model to be used in the virtual space.

volume, an audio source, etc. Every object in Unity can also be scripted using the C# language, meaning that an object can have a user-defined behaviour in the scene. For example, an object can respond to user inputs from a mouse or keyboard or can be translated, rotated and scaled, or have its colour changed, based on conditional events. Object scripting in Unity is key to the modularity of the simulation, especially for the custom representation of PDDL elements.

Scripting can also be applied to the editor window, where users interact with the engine and where it is possible to set the properties of the objects in the scene by using Unity’s user interface. PDSim makes heavy use of all the features provided by Unity, such as the Visual Scripting Language used to create animations and events. As a result, users do not need to learn a new language to develop animations and animation graphs can be modified on the fly without waiting for scripts to be recompiled.

A type in PDSim is represented by a *simulation object*, a structure that shares similar information for all the objects defined in a planning problem. A simulation object is defined by two main components: models and control points. Models are used to visually represent the object type in the virtual world (e.g., block, airport, player, robot, etc.). These can be 3D meshes or 2D textured sprites that can be imported into the Unity editor. A user can add as many models as they like. A collision box that wraps all the models is automatically calculated to be used later in the simulation to detect the interaction with the user inputs and the collisions calculated by the physics engine. Control points are 3D vectors that represent particular points of interest in the object type representation (e.g., the cardinal points of an object, a point that represents the arm position of an agent, etc.).

Figure 10 shows an example of how a simulation object can be composed. The models (highlighted in red) are composed of only one mesh representing a robot rover, and the control points (highlighted in blue) are the 3D vector positions of the front and back of the robot that can be used inside the animations as an anchor point for other objects (e.g., attaching cargo on the front).

If types are specified in the domain definition, then the simulation manager creates simulation object blueprints for



all the leaf types of the type tree that is built when the domain is parsed for the first time. These types are replicated for each object defined in the problem that matches the particular type, using the user configuration of simulation objects, as described above.

A simulation manager is initialised using the Protobuf data from the backend server containing the planning model and the representation of the plan. Every action effect will have an associated list of animation graphs representing the effects of an action. The simulation manager will execute the animations using the attributes in the plan representing the simulation objects involved in the simulation of that action. As the first step in every simulation, the *init* block is animated. Init represents the starting state of a planning problem and is defined by a list of fluents describing the current state of the world. These fluents are represented in the form of *fluent\_name(arguments)* where the arguments are the objects that are present in the environment. The simulation manager will publish events with the corresponding fluent name and objects from the simulation scene that will be used by the visual scripting language to map which animation to execute and the graphical objects to use. The process is then repeated for every action effect in the plan.

## Back-End

PDSim’s backend system is a Python server that communicates with the Unity editor and supports communication between the planning and animation components. Unity tries to connect to the backend server to check if a planning problem has been initialised with the UPL and create a template Unity scene for the visualisation. Planners that can be used by PDSim include FastDownward (Helmert 2006), ENHSP (Scala et al. 2016), Tamer (Valentini, Micheli, and Cimatti 2020), LPG (Gerevini and Serina 2002), Aries (Bit-Monnot 2023) and Pyperplan (Alkhazraji et al. 2020). If either the parsing or planning actions fail, the interface will warn the user of the error.

PDSim’s backend system wraps up the functionality of UPL as the main tool for manipulating and solving planning problems in PDSim. UPL is a Python library provided by the AIPlan4EU project<sup>7</sup> that aims to simplify the use of automated planning tools for AI application development. UPL attempts to standardize aspects of the planning process, making it accessible to users of any level of expertise. In particular, it offers a well-developed PDDL parser and a standard interface for communicating with external planners. Integration with UPL enables PDSim to take advantage of these features and any future updates that UPL may provide.

At the technical level, communication between PDSim’s backend server and Unity is provided by the ZeroMQ networking library,<sup>8</sup> in particular the Python implementation package pyzmq<sup>9</sup> on the server side and the C# implementation netMQ<sup>10</sup> on the Unity side.

<sup>7</sup><https://www.aiplan4eu-project.eu/>

<sup>8</sup><https://zeromq.org/>

<sup>9</sup><https://pypi.org/project/pyzmq/>

<sup>10</sup><https://github.com/zeromq/netmq/>




Figure 11: PDSim for real-world robotics using an HSR robot (Yamamoto et al. 2019). This example represents how recorded semantic sensor data in a ROS project can be visualised as a 3D animation.

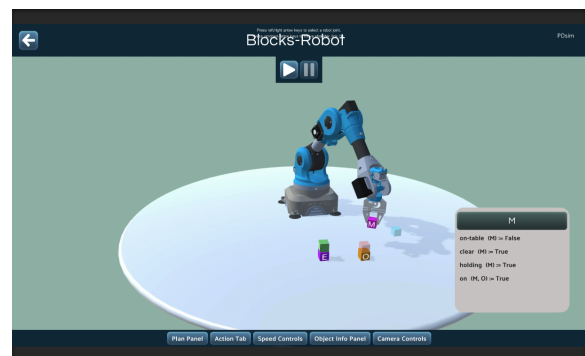


Figure 12: Blocks domain using a simulated Niryo robot<sup>12</sup> and ROS for the arm motion planning.

## Examples

PDSim has been developed and tested using published benchmark domains from the International Planning Competition (IPC),<sup>11</sup> and is currently being used to visualise real-world planning problems. We illustrate the capabilities of PDSim using examples from real-world agricultural and robotics use-cases. A video demonstration of how to set up a simulation with the system is also available.

## Real-World Robotics

PDSim can be used to represent and visualise state changes from real-world robotic scenarios as shown in Figure 11 and 12. In particular, Figure 12 shows a simulation of the blocks planning problem using a robot arm to perform the stack and unstack actions and Figure 11 shows a visualisation of the state changes related to sensors in a smart home. PDSim plays animations related to the robot’s movements between rooms or joint interpolations. This is done by connecting Unity with the Robot Operating System (ROS) and using the specialised library for all kinds of robotic tasks.

<sup>11</sup><https://github.com/potassco/pddl-instances>

<sup>12</sup><https://niryo.com/>



Figure 13: PDSim for an agricultural domain. This example represents a real-world planning problem that involves scheduling trucks and harvesters in fields using a realistic environment that has been digitalized for privacy purposes.

### Agricultural Use-Case

PDSim has been used to visualise a real-world use case involving an agricultural planning problem currently being developed by the Agrotech Valley Forum.<sup>13</sup> The problem involves a scenario as shown in Figure 13 that has been converted into 3D models of roads and fields, with a set of harvesters and vehicles for the transportation of grain into a silo for stocking. Vehicles can only access the fields at particular access points in the map and there is a need for a planning solution to direct the transporting vehicle that follows the harvesting of the fields (Agrotech Valley Forum EV 2023).

### Video Example

Due to the interactive nature of PDSim, we have also created a video to demonstrate the capabilities of the system.<sup>14</sup> This video shows how to start a new simulation, from problem definition to final 3D animation including all the interactions with the Unity front end needed to customise a plan visualisation, as an introduction to PDSim for new users.

### Discussion

In general, PDSim offers a powerful and flexible framework for visualising planning problems using a state-of-the-art graphical engine. More specifically, PDSim aims to fill a gap in current systems that provide plan simulations, by offering users a simplified environment to develop 3D or 2D simulations compared with current approaches that come with the overhead of learning and using an ad hoc scripting language to interact with a custom simulator (Tapia, San Segundo, and Artieda 2015; Chen et al. 2020; Roberts et al. 2021).

PDSim is also designed as a support system for automated planning by providing intuitive tools to interface with a planning solution. Approaches like (Le Bras et al. 2020; Fox, Long, and Magazzeni 2017) also suggest that answering the question of why an action has been successfully executed or has failed, further increases the explainability of a plan. In

<sup>13</sup><https://www.ai4europe.eu/ai-community/organizations/association/agrotech-valley-forum-ev>

<sup>14</sup><https://drive.google.com/file/d/1AH1cYkadRa1ndJp7sxpC2VE0OTEZh0ii/view?usp=sharing>

this context, PDSim provides intuitive hints about possible errors using visual cues, by displaying an interface with the transitions of each action and how they modify the state of a particular object or agent.

However, it is important to reiterate that PDSim is primarily aimed at planning-agnostic users like students. Within this group, as (Chen et al. 2020) indicates, there is a difference between the mental model the user has of the planning problem and the actual implementation. PDDL is often approached as a traditional programming language by beginners, rather than a knowledge definition language. With this in mind, PDSim aims to simplify the learning curve of PDDL by assisting with components that provide information about the state of planning entities in real-time.

### Conclusion and Future Work

This paper presented the structure and operation of PDSim, a simulation system for animating PDDL-based planning domains and plans. In future work, we plan to introduce a more intuitive way to create and modify the knowledge model, using the same visual scripting paradigm and, thus, completely removing the need to know PDDL syntax. This will be internally used together with an in-engine planner that the user can interact with at planning time to change object properties and replan on the fly. Given the close relationship between PDSim and Unity, it will also be possible to use applications such as extended reality (XR) to interact with the plan. Another planned direction for PDSim will also be to include extensions for visualising the current state of an agent’s knowledge and beliefs to support epistemic planning, allowing visualisations to be generated from different agent perspectives. Finally, an evaluation is scheduled to be performed to assess the use of PDSim in an education setting, providing feedback about the overall helpfulness and usefulness of PDSim as a development aid for students learning about automated planning in an introductory AI course.

### References

- Agrotech Valley Forum EV. 2023. <https://www.ai4europe.eu/business-and-industry/case-studies/campaign-planning-silage-maize-harvesting>. Accessed: 2023-12-13.
- Alkhazraji, Y.; Frorath, M.; Grützner, M.; Helmert, M.; Liebetaut, T.; Mattmüller, R.; Ortlieb, M.; Seipp, J.; Sprungenberg, T.; Stahl, P.; and Wülfing, J. 2020. Pyperplan. <https://doi.org/10.5281/zenodo.3700819>.
- Bach, B.; Shi, C.; Heulot, N.; Madhyastha, T.; Grabowski, T.; and Dragicovic, P. 2015. Time curves: Folding time to visualize patterns of temporal evolution in data. *IEEE Transactions on Visualization and Computer Graphics*, 22(1): 559–568.
- Bensalem, S.; Havelund, K.; and Orlandini, A. 2014. Verification and validation meet planning and scheduling. *International Journal on Software Tools for Technology Transfer*, 16: 1–12.
- Bit-Monnot, A. 2023. Experimenting with Lifted Plan-Space Planning as Scheduling: Aries in the 2023 IPC. In *International Planning Competition at ICAPS 2023*.



- Chen, G.; Ding, Y.; Edwards, H.; Chau, C. H.; Hou, S.; Johnson, G.; Sharukh Syed, M.; Tang, H.; Wu, Y.; Yan, Y.; Gil, T.; and Nir, L. 2020. Planimation.
- Cimatti, A.; Micheli, A.; and Roveri, M. 2017. Validating domains and plans for temporal planning via encoding into infinite-state linear temporal logic. In *Proceedings of AAAI*, 3547–3554.
- Craighead, J.; Burke, J.; and Murphy, R. 2008. Using the unity game engine to develop sarge: a case study. In *Proceedings of the IROS Simulation Workshop*.
- De Pellegrin, E. 2020. PDSim: Planning Domain Simulation with the Unity Game Engine. In *Proceedings of the ICAPS Workshop on Knowledge Engineering for Planning and Scheduling (KEPS)*.
- De Pellegrin, E.; and Petrick, R. 2021. Automated Planning and Robotics Simulation with PDSim. In *Proceedings of the ICAPS Workshop on Knowledge Engineering for Planning and Scheduling (KEPS)*.
- De Pellegrin, E.; and Petrick, R. 2022. What Plan? Virtual Plan Visualization with PDSim. In *Proceedings of the ICAPS Workshop on Knowledge Engineering for Planning and Scheduling (KEPS)*.
- De Pellegrin, E.; and Petrick, R. 2023. PDSim: Planning Domain Simulation and Animation with the Unity Game Engine. In *ICAPS 2023 Workshop on Knowledge Engineering for Planning and Scheduling (KEPS)*.
- Echeverria, G.; Lassabe, N.; Degroote, A.; and Lemaignan, S. 2011. Modular open robots simulation engine: Morse. In *Proceedings of ICRA*, 46–51.
- Fort, J.; Hogins, J.; and Davis, N. 2021. Boosting computer vision performance with synthetic data. *Unity Blog*.
- Fox, M.; Long, D.; and Magazzeni, D. 2017. Explainable Planning. In *Proceedings of the IJCAI Workshop on Explainable AI*.
- Ganoni, O.; and Mukundan, R. 2017. A framework for visually realistic multi-robot simulation in natural environment. *arXiv preprint arXiv:1708.01938*.
- Gerevini, A.; and Serina, I. 2002. LPG: A Planner Based on Local Search for Planning Graphs with Action Costs. In *Proceedings of AIPS*, 281–290.
- Ghallab, M.; Nau, D.; and Traverso, P. 2004. *Automated Planning: theory and practice*. Elsevier.
- Green, C.; Platin, J.; Pinol, M.; Trang, A.; and Vij, V. 2020. Robotics simulation in Unity is as easy as 1, 2, 3! *Unity Blog*.
- Haas, J. K. 2014. *A history of the Unity game engine*. Worcester Polytechnic Institute.
- Haslum, P.; Lipovetzky, N.; Magazzeni, D.; and Muise, C. 2019. An introduction to the planning domain definition language. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 13(2): 1–187.
- Helmert, M. 2006. The Fast Downward Planning System. *Journal of Artificial Intelligence Research*, 26: 191–246.
- Hill, A.; Komendantskaya, E.; and Petrick, R. 2020. Proof-Carrying Plans: A Resource Logic for AI Planning. In *International Symposium on Principles and Practice of Declarative Programming (PPDP)*, 1–13.
- Howey, R.; and Long, D. 2003. VAL’s Progress: The Automatic Validation Tool for PDDL2.1 used in the International Planning Competition. In *Proceedings of the ICAPS Workshop on The Competition: Impact, Organization, Evaluation, Benchmarks*.
- Le Bras, P.; Carreno, Y.; Lindsay, A.; Petrick, R.; and Chantler, M. J. 2020. PlanCurves: An Interface for End-Users to Visualise Multi-Agent Temporal Plans. In *Proceedings of the ICAPS Workshop on Knowledge Engineering for Planning and Scheduling (KEPS)*.
- Magnaguagno, M. C.; Fraga Pereira, R.; Móre, M. D.; and Meneguzzi, F. R. 2017. Web planner: A tool to develop classical planning domains and visualize heuristic state-space search. In *ICAPS Workshop on User Interfaces and Scheduling and Planning (UISP)*.
- McDermott, D.; Ghallab, M.; Howe, A.; Knoblock, C.; Ram, A.; Veloso, M.; Weld, D.; and Wilkins, D. 1998. PDDL—The planning domain definition language. Technical Report CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control.
- Micheli, A.; and Bit-Monnot, A. A. 2022. , Unified Planning: A Python Library Making Planning Technology Accessible. In *Proceedings of ICAPS System Demonstration*.
- Muise, C. 2016. Planning.domains. In *Proceedings of ICAPS System Demonstration*.
- Roberts, J. O.; Mastorakis, G.; Lazaruk, B.; Franco, S.; Stokes, A. A.; and Bernardini, S. 2021. vPlanSim: An Open Source Graphical Interface for the Visualisation and Simulation of AI Systems. In *Proceedings of ICAPS*, 486–490.
- Scala, E.; Haslum, P.; Thiébaux, S.; and Ramirez, M. 2016. Interval-based relaxation for general numeric planning. In *Proceedings of ECAI*, 655–663.
- Simpson, R. M.; Kitchin, D. E.; and McCluskey, T. L. 2007. Planning domain definition using GIPO. *The Knowledge Engineering Review*, 22(2): 117–134.
- Tapia, C.; San Segundo, P.; and Artieda, J. 2015. A PDDL-based simulation system. In *Proceedings of the IADIS International Conference Intelligent Systems and Agents*.
- Unity Technologies. 2022. Unity.
- Valentini, A.; Micheli, A.; and Cimatti, A. 2020. Temporal planning with intermediate conditions and effects. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 9975–9982.
- Vaquero, T. S.; Romero, V.; Tonidandel, F.; and Silva, J. R. 2007. itSIMPLE2.0: An Integrated Tool for Designing Planning Domains. In *Proceedings of ICAPS*, 336–343.
- Vodrázka, J.; and Chrpá, L. 2010. Visual design of planning domains. In *Proceedings of ICAPS Workshop on Knowledge Engineering for Planning and Scheduling (KEPS)*, 68–69.
- Yamamoto, T.; Terada, K.; Ochiai, A.; Saito, F.; Asahara, Y.; and Murase, K. 2019. Development of human support robot as the research platform of a domestic mobile manipulator. *ROBOMECH Journal*, 6(1): 1–15.