

Planimation

Gang Chen, Yi Ding, Hugo Edwards, Chong Hin Chau, Sai Hou, Grace Johnson,
Mohammed Sharukh Syed, Haoyuan Tang, Yue Wu, Ye Yan, Gil Tidhar and Nir Lipovetzky

The University of Melbourne
Melbourne, Australia
nir.lipovetzky@unimelb.edu.au

Abstract

Planimation is a modular and extensible open source framework to visualise sequential solutions of planning problems specified in PDDL. We introduce a preliminary declarative PDDL-like animation profile specification, expressive enough to synthesise animations of arbitrary initial states and goals of a benchmark with just a single profile.

Introduction

The adoption of a standard declarative specification of planning tasks through PDDL (McDermott et al. 1998), fostered by International Planning Competitions (McDermott 2000), has boosted the development of solvers by the research community. PDDL specifies a model of the problem, and the planner synthesises a solution which can take the form of a sequence, policy or a tree, depending on the model being solved (Geffner and Bonet 2013). Solutions for classical planning take the form of a sequential plan expressed as text-based keywords identifying the *grounded* actions mapping the initial state to a goal state. Special tools exist to validate the soundness of a plan (Howey, Long, and Fox 2004) given a problem specification in PDDL, and assist on the detection of errors in existing solvers. Other tools have been developed to assist the PDDL modeling process (Vaquero et al. 2007; Barreiro et al. 2012). For instance, itSIMPLE assists the analysis of PDDL by translating state chart diagrams encoded by a modeler into Petri-Nets (Vaquero et al. 2007). Other tools provide insights about *solvers* through search tree visualisations of a given search algorithm (Magnaugno et al. 2017), or visualisations of the internal decision making process of a solver (Chakraborti et al. 2017).

If the PDDL specification is syntactically correct but fails to model the classical planning problem the modeler had in mind, then there is no tool to give feedback and facilitate detecting the source of failure. For example, missing preconditions are prevalent and hard to detect if the only feedback is the name of the actions in a valid plan. *Planimation*, a plan visualizer, intends to close the feedback loop and animate a plan given a PDDL specification, relying on visual cues to help modelers find the sources of mental and model misalignments. Furthermore, the tool diminishes the effort needed to understand and explain the dynamics encoded by PDDL problems, as plans explain themselves visually.

PDDL encodes a transition system declaratively by providing typically a single *domain* file specifying actions and predicates, and a *problem* file specifying objects, the initial state and goal states. Actions changing the valuation of predicates specify the state transition we aim to animate visually. In order to do so, we provide a third PDDL-like *animation* file that specifies a sprite for each object, and the animation behaviour triggered *when a predicate becomes true in a state*. The declarative visual animation language decouples the visualisation engine in the same way PDDL decouples models from solvers. PDDL modelers can extend their problems with a single animation profile and visualize the plans returned by existing solvers.

Animation Profile

The visualization is rendered in a two dimensional canvas. The animation profile allows the specification of the following properties for each `object` or `type`: `x` and `y` as integer coordinates or null value; `color` as a hexadecimal RGB value, pre-specified colour constant or random color; `width` and `height` of the object; `prefabImage` as a base64 string of the object's image¹; `depth` of the object in the canvas; a boolean flag `showname` to specify whether to display the object's name; and an optional `label` string to substitute the default name of the object in the canvas. Given that many domains do not specify all the objects in the problem, we also allow for `custom` object declaration, which adds a visual object in our canvas and permits any predicate to animate it. For example, Blocksworld does not specify objects for the claw and table, which can be added as custom objects.

Predicates defined in the domain file are used by the animation profile to specify the animation behaviour applied to either a specific object, a set of objects, or the objects referred in the predicate parameters. Behaviours are specified by changing any of the object's properties using the `equal` operator in the `:effects` section of the predicate. For example, the predicate `(on ?b1 ?b2)` triggers the effects `(equal (?b1 x) (?b2 x))` and `(equal (?b1 y) (add (?b2 y) (?b2 height) 2))`, changing the `x` value of block `?b1` to the `x` value of `?b2`, and the `y` value of `?b1` to the `y` value of `?b2` plus its `height` and a slack of 2 pixels. Visual effects on predicate objects

¹Base64 image generator: <https://www.base64decode.org/>.

are sufficient to animate state transitions, but they cannot capture fully the visual layout specified in the initial state.

In Blockworld, the initial state contains several *ontable* predicates which can easily specify the y coordinate of a block, but not the x value, as its horizontal position is relative to the number of blocks in the table. For this purpose, we developed special *constraint functions* for visual layouts, inspired by the use of Cassowary toolkit (Badros et al. 1999; Badros, Borning, and Stuckey 2001) in commercial applications such as Mac OS X and IOS 6. The functions supported so far by Planimation are `distributex`, which changes the x value of a set of objects and distributes them along a horizontal plane; `distributey` which changes the y value instead and distributes a set of objects vertically; `distribute_within_objects_vertical` and `distribute_within_objects_horizontal` distribute a set of objects within the bounded visual space of another object (e.g cities); and `distribute_grid_around_point` which distributes a set of objects in a grid like structure. These functions can override the x and y values of an object. A function is invoked using the `assign` keyword in the effects of a predicate, and changes properties of the objects specified in the parameters of the predicate and the function. When a function f is invoked by a predicate p , instead of being applied only over the objects of a single predicate p , it is applied over all the objects of all the true instantiations of p in a state s . That is, the set of objects O of a given function f in a predicate p and state s is defined as $O(f, p, s) = \{o \mid o \in \text{obj}(f), f \in \text{eff}(p), p \in s\}$, where $\text{obj}(f)$ specifies the objects in function f that appear in predicate p . For example, the initial state of Blocksworld is visualized by setting the effect of `(ontable ?b)` as `(assign (?b x) (function distributex (objects ?b) (settings (spacebtwn 40))))`. This function is applied over all the blocks that appear in *ontable* predicates true in the initial state, and distributes them along a horizontal line 40 pixels apart. A second effect `(equal (?b y) 0)` assigns a block to the default y position of the table. The full animation profile for Blocksworld can be viewed at tinyurl.com/yxlt96fp. Other special purpose visual functions are `calculate_label`, which returns the cardinality of a set of objects (e.g number of packages in a truck); `align_middle`, which aligns an object in the center of another object; `apply_smaller`, which changes the width or length of an object; and `draw_line`, which draws a line between two objects.

All these functions have been developed taking into account the diversity of IPC domains. More functions may be needed in the future, but these have been sufficient to animate a variety of domains, namely *Blocksworld*, *Grid*, *Logistics* and *Towers of Hanoi*. All animations are translations in space, scaling, or appear/disappear effects on objects. For more information we refer the reader to the documentation in github.com/planimation.

Planimation

Planimation is a modular and extensible open source framework to visualise sequential solutions of planning problems

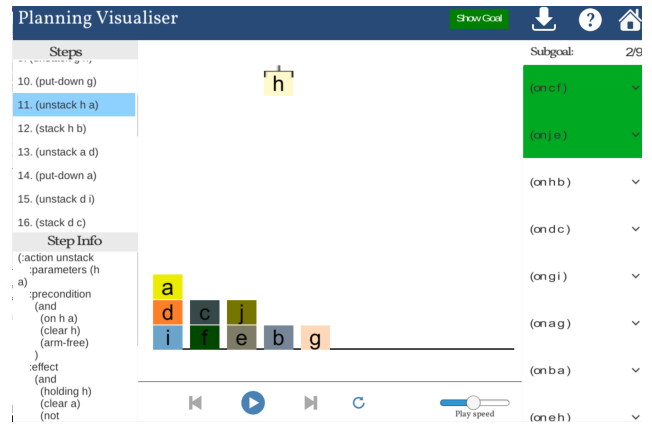


Figure 1: Planimation animating Blocksworld.

specified in PDDL. The framework has two main components, the Unity WebGL *visualiser engine*, and a web API *visualiser solver* written in Python. The solver can either load a saved plan visualisation, or interact with a planner through the API of `solver.planning.domains` (Muisse 2016). Given a domain, problem and animation file, the visualisation solver invokes a planner, retrieves the solution and generates the animations. A visualisation file generator (VFG) can serialise the animation into a json file. Code and documentation for extensions and deployment on dedicated servers are included in github.com/planimation. Planimation runs in the web browser and is accessible through <https://planimation.planning.domains>, as well as through the editor of `planning.domains` via a plugin.

Functionality Planimation UI is shown in Figure 1. It consists of a *steps panel*, which shows the plan and action currently being executed. Any action a is clickable, and sets the visualisation to the resulting state $s' = f(s, a)$. The *Step Info* panel shows the preconditions and effects of the action being executed; the *Animation* panel displays the animation canvas and *darkens* objects if they are in their goal position; the *Control* panel allows animation speed changes and playback controls; the *Subgoal* panel shows all the goal predicates and changes their colour when they are satisfied. Clicking on a subgoal opens a dropdown list with all the steps in which the subgoal is satisfied, jumping to such step if selected. The last panel is the *header*, which contains help information, a button to show the goal state, and a button to download the visualisation in a json VFG format. The file can be used to load planimation without invoking a planner.

Conclusion

Planimation is intended to make solutions more amenable to humans interacting with planners, and assist in the modeling process as well as in the education of AI planning. Such plan visualisations can be part of an explainable AI planning solver (Fox, Long, and Magazzeni 2017) interacting with humans. Planimation has been developed by students as a final year software project, with the vision that the planning community continues its development.

Acknowledgments Naser Soueid and Guangling Yang for preliminary implementations, Gil Tidhar for the supervision of the team, and Christian Muise, Guillem Frances, and Miquel Ramirez for giving feedback on early versions.

References

- Badros, G. J.; Borning, A.; Marriott, K.; and Stuckey, P. 1999. Constraint cascading style sheets for the web. In *Proceedings of the 12th annual ACM symposium on User interface software and technology*, 73–82. ACM.
- Badros, G. J.; Borning, A.; and Stuckey, P. J. 2001. The cassowary linear arithmetic constraint solving algorithm. *ACM Transactions on Computer-Human Interaction (TOCHI)* 8(4):267–306.
- Barreiro, J.; Boyce, M.; Do, M.; Frank, J.; Iatauro, M.; Kichkaylo, T.; Morris, P.; Ong, J.; Remolina, E.; Smith, T.; et al. 2012. Europa: A platform for ai planning, scheduling, constraint programming, and optimization. *4th International Competition on Knowledge Engineering for Planning and Scheduling (ICKEPS)*.
- Chakraborti, T.; Fadnis, K. P.; Talamadupula, K.; Dholakia, M.; Srivastava, B.; Kephart, J. O.; and Bellamy, R. K. 2017. Visualizations for an explainable planning agent. *arXiv preprint arXiv:1709.04517*.
- Fox, M.; Long, D.; and Magazzeni, D. 2017. Explainable planning. *arXiv preprint arXiv:1709.10256*.
- Geffner, H., and Bonet, B. 2013. *A concise introduction to models and methods for automated planning*. Morgan & Claypool Publishers.
- Howey, R.; Long, D.; and Fox, M. 2004. VAL: Automatic plan validation, continuous effects and mixed initiative planning using pddl. In *16th IEEE International Conference on Tools with Artificial Intelligence*, 294–301. IEEE.
- Magnaguagno, M. C.; Pereira, R. F.; Móre, M. D.; and Meneguzzi, F. 2017. Web planner: A tool to develop classical planning domains and visualize heuristic state-space search. In *Proceedings of the Workshop on User Interfaces and Scheduling and Planning, UISP*, 32–38.
- 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, New Haven, CT.
- McDermott, D. 2000. The 1998 AI Planning Systems Competition. *AI Magazine* 21(2):35–56.
- Muise, C. 2016. Planning. domains. *ICAPS system demonstration*.
- Vaquero, T. S.; Romero, V.; Tonidandel, F.; and Silva, J. R. 2007. itsimple 2.0: An integrated tool for designing planning domains. In *ICAPS*, 336–343.