

Solving the Rubik’s Cube with a PDDL Planner

Bharath Muppasani, Vishal Pallagani, Kausik Lakkaraju,
Biplav Srivastava, Forest Agostinelli

AI Institute, University of South Carolina, Columbia, South Carolina, USA

{bharath@email., vishalp@email., kausik@email., biplav.s@, foresta@cse.}sc.edu

Abstract

Rubik’s Cube (RC) is a popular puzzle that is also computationally hard to solve. In this demonstration, we introduce the first PDDL formulation for the 3-dimensional RC and solve it with an off-the-shelf Fast-Downward planner. We also create a plan executor and visualizer to show how the plan achieves the intended goal. Our system has types of two audiences: (a) planning researchers who can explore a hard problem and improve their planning algorithms, and (b) RC learners who want to learn how to solve the puzzle at their own pace and can now modify an initial plan (e.g., manually, using other algorithms) and see their execution.

1 Introduction

As artificial intelligence (AI) continues to solve problems that humans struggle to solve, there is an emerging need for humans to understand these solutions so that we can trust AI, create new educational opportunities, and even discover new knowledge. Many of these problems are path-finding problems. That is, the problem is to find a sequence of actions (a path) to go from any given state to a goal state. AI has been successfully applied to solve the Rubik’s Cube (RC) [Agostinelli *et al.*, 2019; Lakkaraju *et al.*, 2022; Joyner, 2008; Agostinelli *et al.*, 2021] but these methods used opaque learning techniques which are hard for RC learners to benefit from.

While no PDDL encoding of a 3x3x3 RC problem is known to the authors, there is previous work¹ for a 2x2x2 RC setting and is solved with the Fast-Forward planner. Authors in [Büchner *et al.*, 2022] modeled the RC problem in finite domain representation, which enables the common general purpose solvers to be used on the RC problem. Our contributions are: (a) introducing the first PDDL formulation for a 3-size RC; planning researchers can use it to evaluate their planning algorithms, (b) enabling RC learners to use off-the-shelf planners to find custom and optimized way to solve any given RC configuration. Moreover, learning based RC solvers, which have been shown to scale to large instances, can use it as a labeled data generator for training. A demonstration can be seen at the url - here.

¹<https://wu-kan.cn/2019/11/21/Planning-and-Uncertainty/>

2 Background

Rubik’s Cube (RC)

The Rubik’s Cube is a 3-D combination puzzle with colored faces made up of 26 smaller colored pieces linked to a central spindle, with the goal of rotating the blocks until each face of the cube is a single color. To solve the puzzle, one can perform certain actions that correspond to the different faces of the cube. The major actions of a Rubik’s cube are Up(U), Down(D), Right(R), Left(L), Front(F), and Back(B), which define a rotation of 90 degrees in a clockwise direction of the respective face per action. The inverse of these actions corresponds to a 90-degree rotation in the anti-clockwise direction (suffix ‘rev’). One can solve the RC from a scrambled state to the original configuration by performing a set of above mentioned actions.

While sticking to planning terminology of actions, problems and plans, we want to clarify terminology prevalent in RC literature that we will also refer to. A sequence of actions are called *macro-actions* and a collection of macro-actions are called *algorithms* in RC parlance. A solver may employ a strategy for sequencing macro-actions to solve the cube. We use the Fast-Downward AI Planner [Helmert, 2006] to solve the RC problem. Fast Downward is a domain-independent classical planning system based on heuristic search.

3 System Description

We now discuss how an RC can be modelled in PDDL and how the generated plan is linked with a visualizer for providing better understanding to people coming from a non-planning background as well.




Figure 1: Rubik’s cube description to define the domain encoding.

Listing 1: Action L of Rubik’s Cube modeled in PDDL

```

(:action L
:effect (and
;for corner cubelets
(forall(?x ?y ?z)(when (cube1 ?x ?y ?z)
  (and (cube2 ?y ?x ?z))))
(forall(?x ?y ?z)(when (cube3 ?x ?y ?z)
  (and (cube1 ?y ?x ?z))))
(forall(?x ?y ?z)(when (cube4 ?x ?y ?z)
  (and (cube3 ?y ?x ?z))))
(forall(?x ?y ?z)(when (cube2 ?x ?y ?z)
  (and (cube4 ?y ?x ?z))))

;for edge cubelets
(forall(?x ?z)(when (edge13 ?x ?z)
  (and (edge12 ?x ?z))))
(forall(?y ?z)(when (edge34 ?y ?z)
  (and (edge13 ?y ?z))))
(forall(?x ?z)(when (edge24 ?x ?z)
  (and (edge34 ?x ?z))))
(forall(?y ?z)(when (edge12 ?y ?z)
  (and (edge24 ?y ?z))))))

```




Figure 2: Shuffled state of the cube. Solution found with FD planner: U, L. Cost: 2.

three-rubiks-cube npm package² for 3D RC visualization. A random scrambled state of RC represented in the Visualizer is shown in the Figure 2. AI planners are controllable in generating the desired plans. We can specify the search algorithm and the heuristics to the Fast-Downward planner. Each different search algorithm generates different plans to reach the goal state. We employed A* search algorithm in combination with different heuristics which supports conditional-effects in our system and gives plans in minutes.

The system architecture is shown in Figure 3. The users need to upload the domain file and the problem file of RC and select a heuristic of their choice from the dropdown menu. The uploaded domain file and problem file along with the selected heuristic are sent to the API endpoint. The RC visualizer is scrambled to match the initial state from the uploaded problem file. On the back end, the Fast-Downward Planner with A* search along with the selected heuristic evaluates the uploaded problem and generates a solution. This solution is provided to the visualizer to solve the RC. The user may visually follow the actions in the plan file generated to observe the RC being solved step by step. Additional information on the parameters of search time, total time, evaluated states, expanded states, and generated states are also displayed in the front end.

4 Heuristic Evaluation

In our RC solver system, we have incorporated the Fast-Downward AI planner, which offers a range of heuristic implementations for efficient planning. As our RC PDDL employs conditional effects, we have selected a subset of heuristics that support conditional effects in domain modeling. A short description of these heuristics is provided below:

Blind heuristic is based on the idea that the search algorithm doesn’t have any knowledge about the problem domain. It estimates the distance to the goal based on the number of actions needed to reach it, without considering their effects.

Max heuristic estimates the maximum cost of achieving any one of the goals, without considering their interdependencies.

Goal Count heuristic calculates the number of unsatisfied goals in the current state and estimates the cost of satisfying all of them. It works best when there is a small number of goals.

Landmark Cost Partitioning heuristic is based on the idea of breaking down the planning problem into smaller sub-problems, then solving them separately. It uses a technique

²<https://github.com/lab89/three-rubiks-cube>

69 **RC representation in PDDL:** In the PDDL domain, the Ru-
70 bik’s cube problem environment has been defined by assum-
71 ing the cube pieces are in a fixed position, and are named ac-
72 cordingly, as defined in Figure 1. Each action in the RC envi-
73 ronment is defined as the change of colors on these fixed cube
74 pieces. The 3D axis of the cube is considered as three separate
75 parameters X, Y, and Z that specify the position of the color
76 on the cube’s pieces. The three-color cubelet is specified as a
77 predicate with three parameters: X, Y, and Z, which indicate
78 the piece’s colors on three separate axes. The two-color edge
79 piece between the cubelets is specified as a predicate with
80 two parameters denoting the piece’s colors on the two axes.
81 The predicate names define the fixed position of the cubelets
82 and edge pieces that are defined with respect to the different
83 faces of the cube. The representation considered for the cube
84 positions is shown in the Figure 1. One of the actions, ac-
85 tion ‘L’, of RC designed in PDDL of Rubik’s cube from the
86 description provided have been shown in Listing 1.

87 The PDDL for RC is modelled by considering the moves
88 in the RC domain as change of colors on the cube pieces.
89 When the move L is applied to the RC, for example, the left
90 layer is rotated clockwise with respect to the left face. This
91 may be regarded as a clockwise translation of colors from
92 the left layer’s cubelet and edge pieces. When the action L
93 is performed on Figure 1, the colors on the pieces: cube1,
94 cube2, cube4, cube3, are circularly shifted towards bottom.
95 The same applies for the edge pieces. The change of color
96 axis on these pieces is also handled accordingly.

97 **Generating and Visualizing the Plan:** Our system’s PDDL
98 encoded RC solver, in combination with a Visualizer, gen-
99 erates plan actions to solve the problem using the Fast-
100 Downward (FD) AI planner. We have used publicly available

101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145

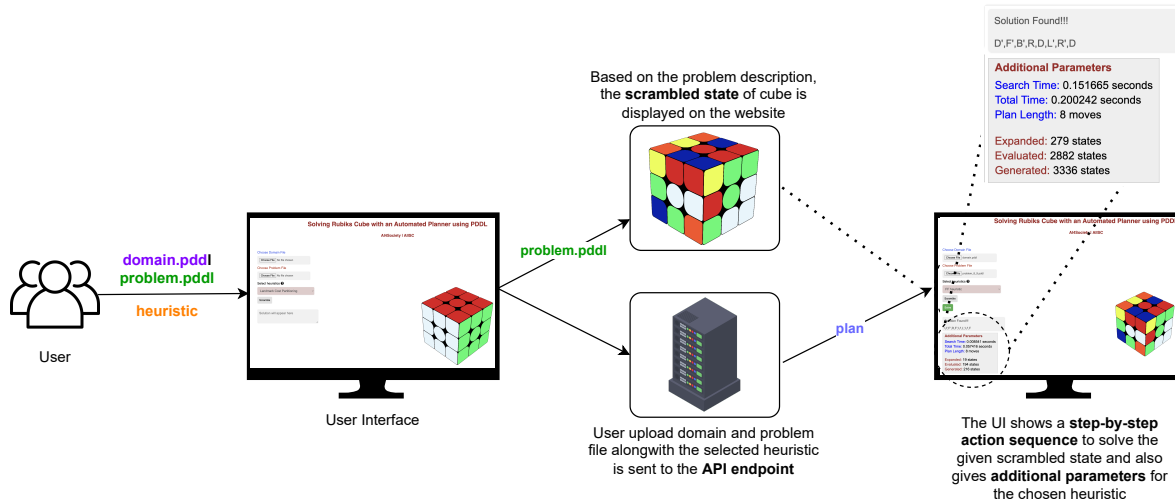


Figure 3: System Architecture.

146 called cost partitioning to distribute the estimated cost of
 147 achieving the goal among these sub-problems. [Karpas and
 148 Domshlak, 2009]

149 *FF* heuristic is based on a simple idea: to achieve a goal,
 150 find the shortest sequence of actions that lead to it. *FF* (Fast
 151 Forward) planner does a forward search from the initial state
 152 to the goal state. [Hoffmann, 2001]

153 *Causal Graph* heuristic constructs a graph representing the
 154 causal relationships between actions and the effects they have
 155 on the state. It estimates the cost of achieving the goal based
 156 on this graph. [Helmert, 2004]

157 *Context-Enhanced Additive* heuristic combines several
 158 sub-heuristics that take into account different aspects of the
 159 problem domain. It estimates the cost of achieving the goal
 160 by adding up the costs of these sub-heuristics and weighting
 161 them based on their relevance to the current state. [Helmert
 162 and Geffner, 2008]

163 In Table 1, we show the performance of these heuristics on
 164 two RC problems (P1 & P2) which are 7 & 8 moves away
 165 from the goal state respectively. Here we compare the heuristics
 166 across time and plan cost metrics. Apart from these metrics,
 167 we also display additional parameters, like the number of
 168 states generated before the solution is found, on our web-
 169 site (as shown in Figure 3). It can be seen that the context-
 170 enhanced additive heuristic is worst performing across these
 171 metrics for both the problems (P1 & P2).

172 5 Conclusion

173 In this work, we have demonstrated the capability of a planner
 174 to solve a complex puzzle, i.e., Rubik’s Cube. For realizing
 175 this, we have created the first PDDL domain for RC. In order
 176 to make the generated plan understandable by people outside
 177 the planning community as well, we have integrated the gener-
 178 ated plan with a visualizer showing step-by-step moves to
 179 achieve a fully solved RC. In the future, we would like to
 180 perform a comparative study of the performance of various
 181 planners and different encodings (PDDL vs. SAS+) to solve
 182 a given RC configuration. Additionally, an empirical study on

Heuristic \ Problem	P1		P2	
	time (s)	cost	time (s)	cost
LM Cost Partitioning	0.3614	7	1.5284	8
FF	0.0777	7	0.1607	8
Goal count	4.8279	7	75.6992	8
Blind	28.1153	7	274.05	8
Max	0.4576	7	3.9793	8
Causal Graph	498.73	21	52.3757	18
Context-enhanced additive	741.146	21	526.99	18

Table 1: Comparison of different heuristics on two RC problems.

183 the performance of abstraction heuristics on the RC modeled
 184 in PDDL, which showed some promising results on SAS+
 185 encoding of RC [Büchner *et al.*, 2022], would be interest-
 186 ing. Integration of a suitable plan validator would be needed
 187 so that human-edited plans can be verified before execution
 188 (currently, VAL [Howey *et al.*, 2004] does not handle condi-
 189 tional effects). This would help us to assist a learner to solve
 190 RC under various constraints such as time or moves.

191 References

- 192 [Agostinelli *et al.*, 2019] Forest Agostinelli, Stephen
 193 McAleer, Alexander Shmakov, and Pierre Baldi. Solving
 194 the Rubik’s cube with deep reinforcement learning and
 195 search. *Nature Mach. Intell.*, 1(8):356–363, 2019.
- 196 [Agostinelli *et al.*, 2021] Forest Agostinelli, Mihir
 197 Mavalankar, Vedant Khandelwal, Hengtao Tang, Dezhi
 198 Wu, Barnett Berry, Biplav Srivastava, Amit Sheth, and
 199 Matthew Irvin. Designing children’s new learning partner:
 200 Collaborative artificial intelligence for learning to solve
 201 the Rubik’s cube. In *Interaction Design and Children*,
 202 pages 610–614, 2021.
- 203 [Büchner *et al.*, 2022] Clemens Büchner, Patrick Ferber,
 204 Jendrik Seipp, and Malte Helmert. A comparison of ab-

205 straction heuristics for rubik’s cube. In *ICAPS 2022 Work-*
206 *shop on Heuristics and Search for Domain-independent*
207 *Planning*, 2022.

208 [Helmert and Geffner, 2008] Malte Helmert and Héctor
209 Geffner. Unifying the causal graph and additive heuristics.
210 In *ICAPS*, pages 140–147, 2008.

211 [Helmert, 2004] Malte Helmert. A planning heuristic based
212 on causal graph analysis. In *ICAPS*, volume 16, pages
213 161–170, 2004.

214 [Helmert, 2006] Malte Helmert. The fast downward plan-
215 ning system. *J. Artif. Int. Res.*, 26(1):191–246, jul 2006.

216 [Hoffmann, 2001] Jörg Hoffmann. Ff: The fast-forward
217 planning system. *AI magazine*, 22(3):57–57, 2001.

218 [Howey *et al.*, 2004] Richard Howey, Derek Long, and
219 Maria Fox. Val: Automatic plan validation, continuous
220 effects and mixed initiative planning using pddl. In *16th*
221 *IEEE International Conference on Tools with Artificial In-*
222 *telligence*, pages 294–301. IEEE, 2004.

223 [Joyner, 2008] David Joyner. *Adventures in group theory:*
224 *Rubik’s Cube, Merlin’s machine, and other mathematical*
225 *toys*. JHU Press, 2008.

226 [Karpas and Domshlak, 2009] Erez Karpas and Carmel
227 Domshlak. Cost-optimal planning with landmarks. In
228 *IJCAI*, pages 1728–1733. Pasadena, CA, 2009.

229 [Lakkaraju *et al.*, 2022] Kaushik Lakkaraju, Thahimum
230 Hassan, Vedant Khandelwal, Prathamjeet Singh, Cassidy
231 Bradley, Ronak Shah, Forest Agostinelli, Biplav Srivas-
232 tava, and Dezhi Wu. ALLURE: A multi-modal guided
233 environment for helping children learn to solve a Rubik’s
234 cube with automatic solving & interactive explanation. In
235 *AAAI*, 2022.