

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 CUBEBENCH: DIAGNOSING INTERACTIVE, LONG-HORIZON SPATIAL REASONING UNDER PARTIAL OBSERVATIONS

Anonymous authors

Paper under double-blind review

## ABSTRACT

Large Language Model (LLM) agents, while proficient in the digital realm, face a significant gap in physical-world deployment due to the challenge of forming and maintaining a robust spatial mental model. We identify three core cognitive challenges hindering this transition: spatial reasoning, long-horizon state tracking via mental simulation, and active exploration under partial observation. To isolate and evaluate these faculties, we introduce **CubeBench**, a novel generative benchmark centered on the Rubik’s Cube. CubeBench uses a three-tiered diagnostic framework that progressively assesses agent capabilities, from foundational state tracking with full symbolic information to active exploration with only partial visual data. Our experiments on leading LLMs reveal critical limitations, including a uniform 0.00% pass rate on all long-horizon tasks, exposing a fundamental failure in long-term planning. We also propose a diagnostic framework to isolate these cognitive bottlenecks by providing external solver tools. By analyzing the failure modes, we provide key insights to guide the development of more physically-grounded intelligent agents.

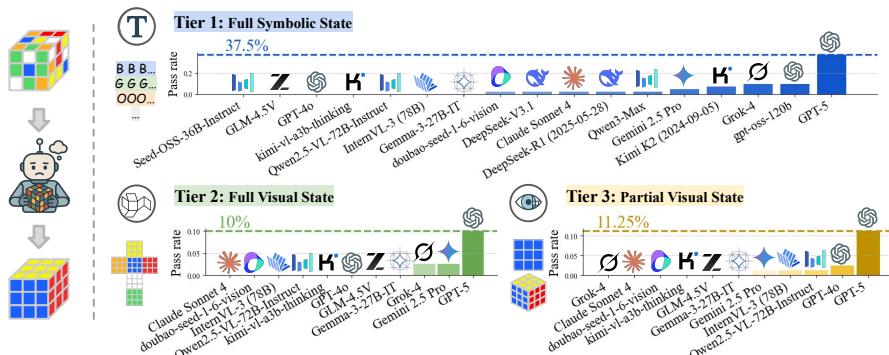


Figure 1: An overview of the performance of leading LLMs on the CubeBench benchmark, broken down by its three diagnostic tiers. **Tier 1 (Full Symbolic State)** tests foundational state tracking using complete symbolic information, where the best average pass rate is only 37.5%. **Tier 2 (Full Visual State)** challenges visual and spatial reasoning by requiring agents to interpret a 2D unfolded map, and **Tier 3 (Partial Visual State)** evaluates active exploration from partial views. Across all tiers, GPT-5 emerges as the top-performing model, though the results highlight a significant performance gap between symbolic and visual reasoning tasks.

## 1 INTRODUCTION

Agents powered by Large Language Models (LLMs) have demonstrated remarkable potential within the digital realm (Gao et al., 2025; Fang et al., 2025). Their proficiency in using tools to navigate websites or write code heralds the dawn of general-purpose AI assistants (Luo et al., 2025; Ma et al., 2025). However, a far grander ambition is to deploy these agents into the physical world. This vision confronts a significant gap: an agent’s success on one-dimensional, symbolic tasks does not readily

translate to effective decision-making in three-dimensional, dynamic environments. The physical world demands more than language comprehension; it requires the ability to form and maintain a robust *spatial mental model* (Johnson-Laird, 1980; 1983).

This gap manifests as critical deficiencies in the core cognitive abilities of current agents, which we show in Fig. 2. We identify three such challenges. The first is **Spatial Reasoning**: physical tasks are inherently three-dimensional, requiring an agent to comprehend an object’s geometry, the relative positions of its components, and the precise consequences of actions in 3D space. The second is **Long-Horizon State Tracking through Mental Simulation**. Unlike digital tasks where state is often externally visible, physical interaction requires an agent to internally maintain and update its world model over long action sequence planning, where even minor errors can accumulate and lead to catastrophic failure. Finally, and most crucially, is the ability for **Exploration and Reasoning under Partial Observation**. The real world rarely provides complete information, so an agent must actively explore its environment to construct a complete mental model from limited views.

To rigorously measure and advance these core capabilities, isolated from the complexities of physical perception, we introduce **CubeBench**, a novel, generative benchmark centered on the Rubik’s Cube. We posit that the cube serves as an ideal laboratory; its deterministic rules and vast state space allow us to conduct controlled experiments that isolate the three core cognitive faculties. To achieve this, CubeBench features a three-tiered diagnostic framework to progressively probe an agent’s capabilities: Tier 1 tests foundational state tracking with complete symbolic information; Tier 2 challenges visual and spatial reasoning by requiring the creation of a 3D model from a 2D unfolded map; and Tier 3 evaluates active exploration using only partial visual information.

Our comprehensive evaluation on CubeBench reveals a staggering performance gap in current LLMs. The results are stark: across all models, **the pass rate on any long-horizon task is a uniform 0.00**, exposing a critical failure in long-term planning and state tracking. Even on short-horizon symbolic tasks, the top-performing LLM, GPT-5<sup>1</sup>, achieves a success rate of just 0.75, merely matching the performance of a traditional Policy Gradient agent and highlighting the difficulty of even basic structured reasoning. Furthermore, our experiments with dense rewards show that while external feedback can provide a local guide on simpler problems, it is insufficient to overcome these core planning deficits. By equipping agents with solver tools, our diagnostic framework successfully pinpoints these failures, isolating long-horizon planning as a primary bottleneck and the inability to reason from partial observations as a more fundamental challenge.

In summary, the primary contributions of this paper are:

- We identify and formalize three core cognitive challenges that impede the deployment of LLM agents into the physical world: spatial reasoning, long-horizon state tracking, and exploration under partial observation.
- We propose CubeBench, a novel, generative benchmark for the controlled evaluation of these cognitive challenges, decoupled from the complexities of visual perception.
- Through extensive experiments on leading LLMs, we reveal their current limitations in forming and utilizing spatial mental models, offering key insights for future development.
- We demonstrate through intervention studies—specifically solver integration and learning from experience—that the identified limitations of base LLMs can be significantly mitigated, pointing toward promising avenues for building more capable agents.

## 2 RELATED WORKS

**Self-evolving Agents.** The paradigm of AI is shifting from static, pre-trained models to dynamic, *self-evolving agents* (Gao et al., 2025; Fang et al., 2025) capable of continual learning and adap-

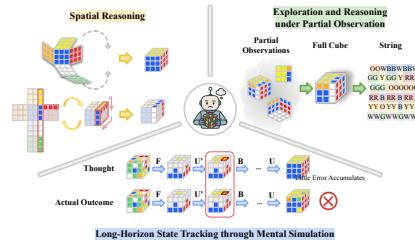


Figure 2: Visualization of the three core cognitive challenges required for **spatial reasoning**.

<sup>1</sup>Accessed via OpenRouter (ID: openai/gpt-5) with unspecified reasoning effort.

108 Table 1: A comparison of agentic benchmarks. Besides the three core cognitive challenges, we also  
 109 evaluate key task characteristics: **Verifiable Outcome Reward**, which assesses if the environment  
 110 operates on fixed, predictable principles rather than subjective or stochastic outcomes such as LLM-  
 111 as-a-judge; **Non-static environment**, which measures if the *state* of the environment changes with  
 112 different agent actions; and whether the task is **Humanly Challenging**, requiring deliberate explo-  
 113 ration for acquiring problem-solving skills beyond simple perception or motor control.

Benchmark Type	Core Cognitive Challenges			Environmental & Task Properties		
	3D Reasoning	Long-Hori. ST. Track	Partial Obs.	Verifiable Outcome Rwd.	Non-static Env.	Humanly Challenging
Search	✗	✓	✗	✓	✗	✓
Code	✗	✓	✓	✓	✓	✓
GUI	✗	✓	✗	✓	✓	✗
Embodied Simulators	✓	✓	✓	✓	✓	✗
Gyms	✗	✓	✓	✓	✓	✓
ARC-AGI-3	✗	✓	✗	✓	✓	✓
MINDCUBE	✓	✗	✗	✓	✗	✗
<b>CubeBench (Ours)</b>	✓	✓	✓	✓	✓	✓

120 ✓ for being explicitly designed to test the capability. ✓ for being partially tested in the benchmark. ✗ for benchmarks not primarily focusing on this capability.

121 tation from experience (Wang et al., 2024a;b; Luo et al., 2025; Zhang et al., 2025b; Hu et al.,  
 122 2024a; 2025; Liang et al., 2024a; Ma et al., 2025). Unlike foundational agents with fixed capa-  
 123 bilities, self-evolving agents can autonomously modify their own components—including memory  
 124 (Zhang et al., 2024a; Zhou et al., 2024; Liang et al., 2024b; Xu et al., 2025b; Zhao et al., 2024a;  
 125 Chhikara et al., 2025; Guan et al., 2024; Yu et al., 2025), tools (Qiu et al., 2025; Haque et al., 2025;  
 126 Zheng et al., 2025; Zhao et al., 2024b; Qu et al., 2025; Wang et al., 2025), and architecture (Zhuang  
 127 et al., 2025; Zhang et al., 2025c; Sapkota et al., 2025) — in response to environmental interaction.  
 128 As these agents evolve to tackle the physical world, a fundamental shift in their evaluation is re-  
 129 quired—moving beyond traditional static assessments to benchmarks that can rigorously measure  
 130 the acquisition and application of *spatial intelligence*.

131 **Benchmarks for Self-evolving Agents.** Existing benchmarks (Chan et al., 2024; Chen et al., 2024b;  
 132 Wei et al., 2025; Levy et al., 2024; Wu et al., 2025; Mialon et al., 2023; Liu et al., 2023; Chen et al.,  
 133 2025a; Zhu et al., 2025; Hu et al., 2024b), however, are not designed for these dynamics. As shown  
 134 in Table 1, different categories of benchmarks test these cognitive skills to varying degrees, but none  
 135 provides a focused, isolated evaluation. Digital environments for Search and GUI interaction (Xie  
 136 et al., 2024; Zhang et al., 2025a; Levy et al., 2024; Wu et al., 2025; Zhou et al., 2023; Deng et al.,  
 137 2023; Mialon et al., 2023; Wei et al., 2025; Phan et al., 2025), for instance, are primarily 2D and  
 138 feature explicit states, thus not addressing 3D spatial reasoning. While Code and Gym environments  
 139 (Hu et al., 2024b; Jimenez et al., 2023; Chan et al., 2024; Chen et al., 2025b; Aleithan et al., 2024;  
 140 Yang et al., 2024; Xu et al., 2024; Su et al., 2025; Tassa et al., 2018; Yu et al., 2020; Rajeswaran  
 141 et al., 2017) require long-horizon state tracking, they do not involve the complex 3D geometric  
 142 understanding that is crucial for physical-world tasks. Embodied simulators (Gao et al., 2024; Yang  
 143 et al., 2025b; Li et al., 2024; Savva et al., 2019b; Shridhar et al., 2021; Kolve et al., 2017) do engage  
 144 all three faculties but inherently couple them with complex visual perception, making it difficult to  
 145 isolate cognitive failures. While recent work like MindCube (Liu et al., 2024) evaluates reasoning  
 146 on static 3D scenes, our work introduces the challenge of updating a spatial model through long-  
 147 horizon, state-altering interaction.

148 In this work, we develop CubeBench, which is specifically designed to fill this gap by decoupling  
 149 perception from reasoning. Its deterministic, rule-based nature makes it an ideal suite for studying  
 150 an agent’s evolution; when an agent fails, the cause can be precisely attributed to a failure in its  
 151 internal spatial model or its long-horizon planning, as shown in Sec. 4.3. Furthermore, CubeBench’s  
 152 generative nature allows for the creation of a virtually infinite curriculum of tasks with fine-grained  
 153 difficulty, enabling the rigorous evaluation of an agent’s ability to learn and adapt over time—a  
 154 cornerstone of assessing true self-evolution (Gao et al., 2025).

### 155 3 THE CUBE BENCH BENCHMARK

#### 156 3.1 TASK DEFINITION

157 We formalize the Rubik’s Cube challenge as a Partially Observable Markov Decision Process  
 158 (POMDP), providing a structured framework to analyze agent behavior. A POMDP is defined by  
 159 a tuple  $(S, A, T, R, \Omega, O)$ , where  $S$  is a set of states,  $A$  is a set of actions,  $T$  is the state transition  
 160 function,  $R$  is the reward function,  $\Omega$  is a set of observations, and  $O$  is the observation function. In

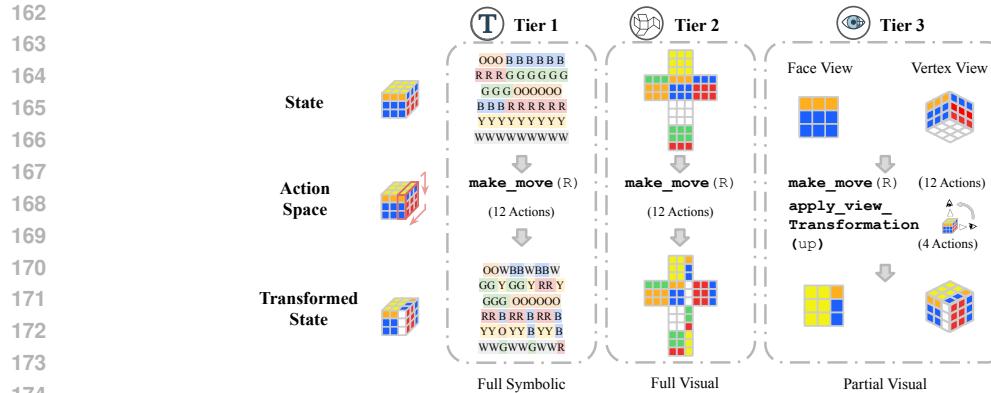


Figure 3: **Illustration on the three-tiered task of CubeBench.** Tier 1 (Full Symbolic State) provides the agent with complete state information in a string format, which makes the problem a fully observable MDP. Tier 2 (Full Visual State) presents the full state as a 2D unfolded map, which challenges the agent’s visual thinking. Tier 3 (Partial Visual State) provides only a partial view of the cube (Face view or Vertex view), which requires the agent to explore the environment to gather the full state information.

this context, the agent’s goal is to learn a policy  $\pi(a|o)$  that selects an action  $a \in A$  given an observation  $o \in \Omega$  to maximize the expected cumulative reward. We now define each of these components within the CubeBench environment.

### 3.1.1 STATE SPACE

The state space  $S$  encompasses all possible configurations of the 3x3x3 Rubik’s Cube. The internal state of the cube,  $s \in S$ , is deterministically represented by a data structure that tracks the color of the 54 individual facelets (stickers). This symbolic representation is unambiguous and allows for perfect state tracking within the simulation. The Rubik’s Cube is a classic example of a system governed by the principles of group theory. Each move corresponds to a permutation of the cube’s facelets, and the set of all possible move sequences forms a mathematical group. This deterministic, non-stochastic nature makes it an ideal environment for isolating an agent’s reasoning and planning capabilities from the complexities of physical uncertainty. The state space is vast, containing over 43 quintillion ( $4.3 \times 10^{19}$ ) unique configurations, yet it is finite and structured. This combination of immense scale and deterministic rules makes it a compelling microcosm for studying autonomous problem-solving on tasks that are too large for naive search but are perfectly predictable.

### 3.1.2 OBSERVATION SPACE

The observation space  $\Omega$  is defined by the observation function  $O(s)$ , which maps the true internal state  $s$  to an observation  $o$  that is presented to the agent. As shown in Fig. 3, CubeBench features a three-tiered observation space, where each tier presents the state information in a different modality, posing distinct perceptual challenges.

**Tier 1: Full Symbolic State.** The observation is a 54-character string that symbolically represents the complete state of the cube. Each character corresponds to the color of a single facelet (e.g., ‘W’ for White, ‘R’ for Red, ‘B’ for Blue, ‘O’ for Orange, ‘G’ for Green, ‘Y’ for Yellow). In this tier, the observation function provides the full state information as a 54-character structured string, making the problem a fully observable MDP.

**Tier 2: Full Visual State.** The observation is a single image depicting the cube’s complete 2D unfolded map, which visually presents all 54 facelets in a planar layout. This tier specifically challenges an agent’s visual reasoning capabilities, requiring it to mentally *fold* the 2D layout into a coherent 3D spatial model to understand the adjacency of faces that are not contiguous in the planar representation.

**Tier 3: Partial Visual State.** The observation is a single image of a partial view of the cube. This can be either an image of a single face (*face view*) or an image from a corner’s perspective showing

216 three adjacent faces (*vertex view*). In this tier, the observation function provides incomplete state  
 217 information, thus formulating the task as a true POMDP.  
 218

219 **3.1.3 ACTION SPACE**  
 220

221 The action space  $A$  consists of the set of discrete, deterministic commands an agent can execute to  
 222 interact with the environment.

223 **State Transition Actions:** The primary action for rotating the cube, which implements the environ-  
 224 ment’s transition function  $T(s, a)$ . A `make_move` command accepts one of 12 standard Singmaster  
 225 notation inputs corresponding to a 90-degree rotation of a face: F (Front), B (Back), L (Left), R  
 226 (Right), U (Up), D (Down), and their counter-clockwise prime versions (F', B', etc.).

227 **Observation-Altering Actions:** An action exclusive to Tier 3 that allows the agent to change its  
 228 observational viewpoint (i.e., up, down, left, right) without altering the cube’s underlying state  $s$ .  
 229 This is the primary mechanism for exploration in the partially observable setting.  
 230

231 **3.1.4 REWARD FUNCTION**  
 232

233 The agent’s objective is to reach the solved state. We define two types of reward signals within  
 234 CubeBench to facilitate and evaluate this process.

235 **Sparse Terminal Reward:** The primary success metric is a sparse, binary reward. The agent re-  
 236 ceives a reward of  $R = 1$  upon entering the terminal solved state (i.e., all stickers on six faces  
 237 are matched), and  $R = 0$  for all other state transitions. The agent’s goal is to find a policy that  
 238 maximizes the probability of achieving this terminal reward within the given constraints.

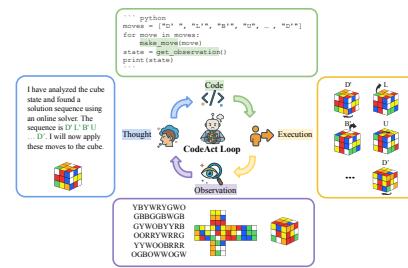
239 **Dense Progressive Reward:** To potentially guide the agent’s search process, we also implement an  
 240 optional *dense* reward mechanism. Unlike a state-value function, our dense rewards are calculated  
 241 as the *change* in a given metric before and after a state transition action. Specifically, the reward  
 242  $R_t$  for taking state transition action  $a_t$  in state  $s_t$  to reach state  $s_{t+1}$  is defined as the difference in a  
 243 metric function  $\phi(s)$ :

$$R_t = \phi(s_{t+1}) - \phi(s_t)$$

244 We implemented and tested three different metric functions ( $\phi$ ) to explore how the conceptual gran-  
 245 ularity of the feedback affects agent performance. (1) **Sticker Metric** ( $\phi_{sticker}$ ): This function  
 246 quantifies the total number of individual facelets (stickers) that are in their correct home positions.  
 247 The score  $\phi_{sticker}(s)$  ranges from 9 for a highly scrambled cube to 54 for the solved state. This  
 248 provides a fine-grained, low-level signal of progress. (2) **Face Metric** ( $\phi_{face}$ ): This function counts  
 249 the number of fully solved faces, where all 9 stickers on a face are correct. It provides a high-level,  
 250 more conceptually grounded signal that is sparser than the sticker metric. (3) **Heuristic Metric**  
 251 ( $\phi_{heuristic}$ ): This function uses an algorithmic heuristic from a common solving method to estimate  
 252 the distance to the goal state. It is designed to provide a more informed, albeit abstract, numerical  
 253 signal, which we explain in detail in Sec. J. As a default setting, we also include a `no_reward`  
 254 condition where  $R_t = 0$  for all transitions.  
 255

256 **3.2 TASK EVALUATION AND GENERATION**  
 257

258 **Agent Interaction Protocol.** The agent’s interaction  
 259 with the environment follows the ReAct paradigm (Yao  
 260 et al. (2022)), structured into a sequence of decision-  
 261 making steps. As shown in Fig. 4, we define a single *step*  
 262 as a complete Thought–Code–Observation block.  
 263 Within each step, the agent first generates its reasoning  
 264 (Thought), then writes and executes code to interact  
 265 with the environment (Code), and finally receives the  
 266 output of that code as feedback (Observation) for its  
 267 next cycle. Each experimental run is subject to a maxi-  
 268 mum of 20 steps and a timeout of 30 minutes to ensure  
 269 fair comparison. Note that in each step, the agent could  
 write code to make more than one move.



259  
 260  
 261  
 262  
 263  
 264  
 265  
 266  
 267  
 268  
 269  
 270  
 271  
 272  
 273  
 274  
 275  
 276  
 277  
 278  
 279  
 280  
 281  
 282  
 283  
 284  
 285  
 286  
 287  
 288  
 289  
 290  
 291  
 292  
 293  
 294  
 295  
 296  
 297  
 298  
 299  
 300  
 301  
 302  
 303  
 304  
 305  
 306  
 307  
 308  
 309  
 310  
 311  
 312  
 313  
 314  
 315  
 316  
 317  
 318  
 319  
 320  
 321  
 322  
 323  
 324  
 325  
 326  
 327  
 328  
 329  
 330  
 331  
 332  
 333  
 334  
 335  
 336  
 337  
 338  
 339  
 340  
 341  
 342  
 343  
 344  
 345  
 346  
 347  
 348  
 349  
 350  
 351  
 352  
 353  
 354  
 355  
 356  
 357  
 358  
 359  
 360  
 361  
 362  
 363  
 364  
 365  
 366  
 367  
 368  
 369  
 370  
 371  
 372  
 373  
 374  
 375  
 376  
 377  
 378  
 379  
 380  
 381  
 382  
 383  
 384  
 385  
 386  
 387  
 388  
 389  
 390  
 391  
 392  
 393  
 394  
 395  
 396  
 397  
 398  
 399  
 400  
 401  
 402  
 403  
 404  
 405  
 406  
 407  
 408  
 409  
 410  
 411  
 412  
 413  
 414  
 415  
 416  
 417  
 418  
 419  
 420  
 421  
 422  
 423  
 424  
 425  
 426  
 427  
 428  
 429  
 430  
 431  
 432  
 433  
 434  
 435  
 436  
 437  
 438  
 439  
 440  
 441  
 442  
 443  
 444  
 445  
 446  
 447  
 448  
 449  
 450  
 451  
 452  
 453  
 454  
 455  
 456  
 457  
 458  
 459  
 460  
 461  
 462  
 463  
 464  
 465  
 466  
 467  
 468  
 469  
 470  
 471  
 472  
 473  
 474  
 475  
 476  
 477  
 478  
 479  
 480  
 481  
 482  
 483  
 484  
 485  
 486  
 487  
 488  
 489  
 490  
 491  
 492  
 493  
 494  
 495  
 496  
 497  
 498  
 499  
 500  
 501  
 502  
 503  
 504  
 505  
 506  
 507  
 508  
 509  
 510  
 511  
 512  
 513  
 514  
 515  
 516  
 517  
 518  
 519  
 520  
 521  
 522  
 523  
 524  
 525  
 526  
 527  
 528  
 529  
 530  
 531  
 532  
 533  
 534  
 535  
 536  
 537  
 538  
 539  
 540  
 541  
 542  
 543  
 544  
 545  
 546  
 547  
 548  
 549  
 550  
 551  
 552  
 553  
 554  
 555  
 556  
 557  
 558  
 559  
 560  
 561  
 562  
 563  
 564  
 565  
 566  
 567  
 568  
 569  
 570  
 571  
 572  
 573  
 574  
 575  
 576  
 577  
 578  
 579  
 580  
 581  
 582  
 583  
 584  
 585  
 586  
 587  
 588  
 589  
 590  
 591  
 592  
 593  
 594  
 595  
 596  
 597  
 598  
 599  
 600  
 601  
 602  
 603  
 604  
 605  
 606  
 607  
 608  
 609  
 610  
 611  
 612  
 613  
 614  
 615  
 616  
 617  
 618  
 619  
 620  
 621  
 622  
 623  
 624  
 625  
 626  
 627  
 628  
 629  
 630  
 631  
 632  
 633  
 634  
 635  
 636  
 637  
 638  
 639  
 640  
 641  
 642  
 643  
 644  
 645  
 646  
 647  
 648  
 649  
 650  
 651  
 652  
 653  
 654  
 655  
 656  
 657  
 658  
 659  
 660  
 661  
 662  
 663  
 664  
 665  
 666  
 667  
 668  
 669  
 670  
 671  
 672  
 673  
 674  
 675  
 676  
 677  
 678  
 679  
 680  
 681  
 682  
 683  
 684  
 685  
 686  
 687  
 688  
 689  
 690  
 691  
 692  
 693  
 694  
 695  
 696  
 697  
 698  
 699  
 700  
 701  
 702  
 703  
 704  
 705  
 706  
 707  
 708  
 709  
 710  
 711  
 712  
 713  
 714  
 715  
 716  
 717  
 718  
 719  
 720  
 721  
 722  
 723  
 724  
 725  
 726  
 727  
 728  
 729  
 730  
 731  
 732  
 733  
 734  
 735  
 736  
 737  
 738  
 739  
 740  
 741  
 742  
 743  
 744  
 745  
 746  
 747  
 748  
 749  
 750  
 751  
 752  
 753  
 754  
 755  
 756  
 757  
 758  
 759  
 760  
 761  
 762  
 763  
 764  
 765  
 766  
 767  
 768  
 769  
 770  
 771  
 772  
 773  
 774  
 775  
 776  
 777  
 778  
 779  
 780  
 781  
 782  
 783  
 784  
 785  
 786  
 787  
 788  
 789  
 790  
 791  
 792  
 793  
 794  
 795  
 796  
 797  
 798  
 799  
 800  
 801  
 802  
 803  
 804  
 805  
 806  
 807  
 808  
 809  
 810  
 811  
 812  
 813  
 814  
 815  
 816  
 817  
 818  
 819  
 820  
 821  
 822  
 823  
 824  
 825  
 826  
 827  
 828  
 829  
 830  
 831  
 832  
 833  
 834  
 835  
 836  
 837  
 838  
 839  
 840  
 841  
 842  
 843  
 844  
 845  
 846  
 847  
 848  
 849  
 850  
 851  
 852  
 853  
 854  
 855  
 856  
 857  
 858  
 859  
 860  
 861  
 862  
 863  
 864  
 865  
 866  
 867  
 868  
 869  
 870  
 871  
 872  
 873  
 874  
 875  
 876  
 877  
 878  
 879  
 880  
 881  
 882  
 883  
 884  
 885  
 886  
 887  
 888  
 889  
 890  
 891  
 892  
 893  
 894  
 895  
 896  
 897  
 898  
 899  
 900  
 901  
 902  
 903  
 904  
 905  
 906  
 907  
 908  
 909  
 910  
 911  
 912  
 913  
 914  
 915  
 916  
 917  
 918  
 919  
 920  
 921  
 922  
 923  
 924  
 925  
 926  
 927  
 928  
 929  
 930  
 931  
 932  
 933  
 934  
 935  
 936  
 937  
 938  
 939  
 940  
 941  
 942  
 943  
 944  
 945  
 946  
 947  
 948  
 949  
 950  
 951  
 952  
 953  
 954  
 955  
 956  
 957  
 958  
 959  
 960  
 961  
 962  
 963  
 964  
 965  
 966  
 967  
 968  
 969  
 970  
 971  
 972  
 973  
 974  
 975  
 976  
 977  
 978  
 979  
 980  
 981  
 982  
 983  
 984  
 985  
 986  
 987  
 988  
 989  
 990  
 991  
 992  
 993  
 994  
 995  
 996  
 997  
 998  
 999  
 1000  
 1001  
 1002  
 1003  
 1004  
 1005  
 1006  
 1007  
 1008  
 1009  
 1010  
 1011  
 1012  
 1013  
 1014  
 1015  
 1016  
 1017  
 1018  
 1019  
 1020  
 1021  
 1022  
 1023  
 1024  
 1025  
 1026  
 1027  
 1028  
 1029  
 1030  
 1031  
 1032  
 1033  
 1034  
 1035  
 1036  
 1037  
 1038  
 1039  
 1040  
 1041  
 1042  
 1043  
 1044  
 1045  
 1046  
 1047  
 1048  
 1049  
 1050  
 1051  
 1052  
 1053  
 1054  
 1055  
 1056  
 1057  
 1058  
 1059  
 1060  
 1061  
 1062  
 1063  
 1064  
 1065  
 1066  
 1067  
 1068  
 1069  
 1070  
 1071  
 1072  
 1073  
 1074  
 1075  
 1076  
 1077  
 1078  
 1079  
 1080  
 1081  
 1082  
 1083  
 1084  
 1085  
 1086  
 1087  
 1088  
 1089  
 1090  
 1091  
 1092  
 1093  
 1094  
 1095  
 1096  
 1097  
 1098  
 1099  
 1100  
 1101  
 1102  
 1103  
 1104  
 1105  
 1106  
 1107  
 1108  
 1109  
 1110  
 1111  
 1112  
 1113  
 1114  
 1115  
 1116  
 1117  
 1118  
 1119  
 1120  
 1121  
 1122  
 1123  
 1124  
 1125  
 1126  
 1127  
 1128  
 1129  
 1130  
 1131  
 1132  
 1133  
 1134  
 1135  
 1136  
 1137  
 1138  
 1139  
 1140  
 1141  
 1142  
 1143  
 1144  
 1145  
 1146  
 1147  
 1148  
 1149  
 1150  
 1151  
 1152  
 1153  
 1154  
 1155  
 1156  
 1157  
 1158  
 1159  
 1160  
 1161  
 1162  
 1163  
 1164  
 1165  
 1166  
 1167  
 1168  
 1169  
 1170  
 1171  
 1172  
 1173  
 1174  
 1175  
 1176  
 1177  
 1178  
 1179  
 1180  
 1181  
 1182  
 1183  
 1184  
 1185  
 1186  
 1187  
 1188  
 1189  
 1190  
 1191  
 1192  
 1193  
 1194  
 1195  
 1196  
 1197  
 1198  
 1199  
 1200  
 1201  
 1202  
 1203  
 1204  
 1205  
 1206  
 1207  
 1208  
 1209  
 1210  
 1211  
 1212  
 1213  
 1214  
 1215  
 1216  
 1217  
 1218  
 1219  
 1220  
 1221  
 1222  
 1223  
 1224  
 1225  
 1226  
 1227  
 1228  
 1229  
 1230  
 1231  
 1232  
 1233  
 1234  
 1235  
 1236  
 1237  
 1238  
 1239  
 1240  
 1241  
 1242  
 1243  
 1244  
 1245  
 1246  
 1247  
 1248  
 1249  
 1250  
 1251  
 1252  
 1253  
 1254  
 1255  
 1256  
 1257  
 1258  
 1259  
 1260  
 1261  
 1262  
 1263  
 1264  
 1265  
 1266  
 1267  
 1268  
 1269  
 1270  
 1271  
 1272  
 1273  
 1274  
 1275  
 1276  
 1277  
 1278  
 1279  
 1280  
 1281  
 1282  
 1283  
 1284  
 1285  
 1286  
 1287  
 1288  
 1289  
 1290  
 1291  
 1292  
 1293  
 1294  
 1295  
 1296  
 1297  
 1298  
 1299  
 1300  
 1301  
 1302  
 1303  
 1304  
 1305  
 1306  
 1307  
 1308  
 1309  
 1310  
 1311  
 1312  
 1313  
 1314  
 1315  
 1316  
 1317  
 1318  
 1319  
 1320  
 1321  
 1322  
 1323  
 1324  
 1325  
 1326  
 1327  
 1328  
 1329  
 1330  
 1331  
 1332  
 1333  
 1334  
 1335  
 1336  
 1337  
 1338  
 1339  
 1340  
 1341  
 1342  
 1343  
 1344  
 1345  
 1346  
 1347  
 1348  
 1349  
 1350  
 1351  
 1352  
 1353  
 1354  
 1355  
 1356  
 1357  
 1358  
 1359  
 1360  
 1361  
 1362  
 1363  
 1364  
 1365  
 1366  
 1367  
 1368  
 1369  
 1370  
 1371  
 1372  
 1373  
 1374  
 1375  
 1376  
 1377  
 1378  
 1379  
 1380  
 1381  
 1382  
 1383  
 1384  
 1385  
 1386  
 1387  
 1388  
 1389  
 1390  
 1391  
 1392  
 1393  
 1394  
 1395  
 1396  
 1397  
 1398  
 1399  
 1400  
 1401  
 1402  
 1403  
 1404  
 1405  
 1406  
 1407  
 1408  
 1409  
 1410  
 1411  
 1412  
 1413  
 1414  
 1415  
 1416  
 1417  
 1418  
 1419  
 1420  
 1421  
 1422  
 1423  
 1424  
 1425  
 1426  
 1427  
 1428  
 1429  
 1430  
 1431  
 1432  
 1433  
 1434  
 1435  
 1436  
 1437  
 1438  
 1439  
 1440  
 1441  
 1442  
 1443  
 1444  
 1445  
 1446  
 1447  
 1448  
 1449  
 1450  
 1451  
 1452  
 1453  
 1454  
 1455  
 1456  
 1457  
 1458  
 1459  
 1460  
 1461  
 1462  
 1463  
 1464  
 1465  
 1466  
 1467  
 1468  
 1469  
 1470  
 1471  
 1472  
 1473  
 1474  
 1475  
 1476  
 1477  
 1478  
 1479  
 1480  
 1481  
 1482  
 1483  
 1484  
 1485  
 1486  
 1487  
 1488  
 1489  
 1490  
 1491  
 1492  
 1493  
 1494  
 1495  
 1496  
 1497  
 1498  
 1499  
 1500  
 1501  
 1502  
 1503  
 1504  
 1505  
 1506  
 1507  
 1508  
 1509  
 1510  
 1511  
 1512  
 1513  
 1514  
 1515  
 1516  
 1517  
 1518  
 1519  
 15

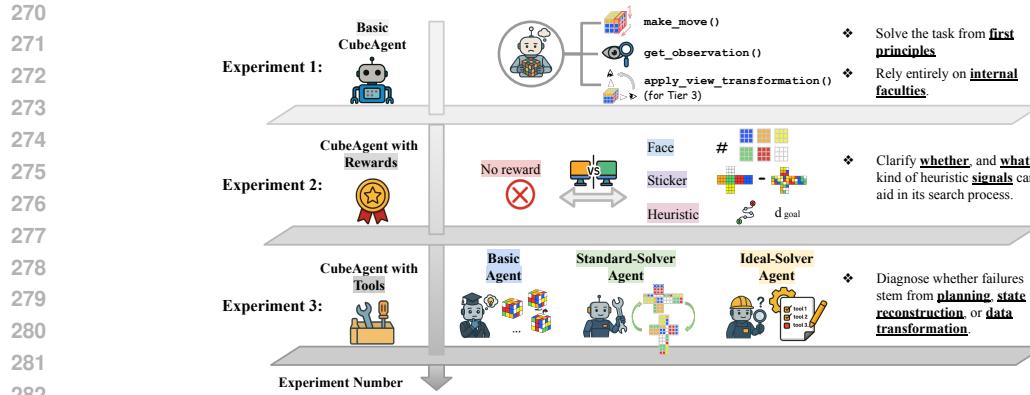


Figure 5: **Visualization of our three-part diagnostic framework for systematically evaluating LLM agents.** To answer **Q1**, we test a basic agent with only fundamental interaction tools to establish its baseline capabilities from first principles. For **Q2**, we augment the agent with various dense reward signals to determine if external feedback can effectively guide its search process. Finally, to address **Q3**, we deploy agents with different levels of tool support to diagnose whether failures originate from high-level planning, state reconstruction, or procedural data transformation.

**Evaluation Metrics.** Our primary metrics are designed to measure both success and effort. The *Pass Rate* is the fraction of test cases successfully solved within the execution constraints, serving as the primary measure of an agent’s capability. To quantify search effort, we use the *Number of make\_move calls (#MM)* as a proxy. We report this metric aggregated in three ways: the average over all normally terminated runs, the average over only successful runs, and the maximum count observed across all normally terminated runs.

**Task Generation and Difficulty Scaling.** We define task difficulty based on the *optimal* number of moves required to solve a given cube configuration, a metric we refer to as the state’s *depth*. A state’s depth serves as a robust proxy for its complexity; solving high-depth states is infeasible through random exploration and necessitates a coherent strategy. To generate our test cases, we employ a provably optimal solver (see Appendix K). For a target depth  $d$ , we generate scrambled states and confirm their optimality by verifying that a solution of length  $d$  exists, but no solution of length  $d - 1$  can be found. This guarantees the true depth is precisely  $d$ . The detail of this process is described in the appendix. To analyze agent performance across varying complexities, we group these cases into two distinct categories: *Short-Horizon* tasks, comprising states with depths of 1, 2, 3, and 4, and *Long-Horizon* tasks, which include the more challenging depths of 8, 12, 16, and 20. The configuration of the generated test split is described in detail in Sec.G.

## 4 DIAGNOSING LLM AGENT CAPABILITIES ON CUBE BENCH

In this section, we introduce our systematic framework for evaluating Large Language Model (LLM) agents on the CubeBench benchmark. As shown in Fig. 5, the evaluation process is designed as a three-part diagnostic, structured around three central research questions that aim to progressively uncover the cognitive strengths and weaknesses of current agents:

### Questions

- **Q1:** What are the baseline capabilities and limitations of current LLM agents when trying to solve the typical cube problem in an unaided setting?
- **Q2:** Can the introduction of dense reward signals effectively guide an agent’s context-based reasoning process and enhance its performance on these complex spatial tasks?
- **Q3:** How can we design a diagnostic evaluation to isolate the impact of each core cognitive challenge, thereby identifying the primary bottlenecks for agent failure—is it high-level planning, state reconstruction from partial perception, or spatial reasoning?

Table 2: Baseline performance across modalities and horizons on CubeBench. Top row: metric groups (**Pass rate**, **#MM**, where #MM is the average number of `make_move` calls); second row: observation modalities; **third row: task horizons (Short = S, depths 1–4; Long = L, depths 8, 12, 16, 20)**. **Tier 3** denotes the hardest split, evaluated under two projections: *Face View* and *Vertex View*. Blue shading denotes open-source models, and pink denotes proprietary models. For each metric column, we shade the top-3 entries (**red** = 1st, **orange** = 2nd, **yellow** = 3rd). We also train an MLP with policy gradient on the Full Symbolic setting; details are in Sec. H. “**—**” : Model does not support visual inputs.

Model	Pass rate								#MM							
	Full Symbolic		Full Visual		Face view		Vertex view		Full Symbolic		Full Visual		Face view		Vertex view	
	S	L	S	L	S	L	S	L	S	L	S	L	S	L	S	L
GPT-5	<b>0.75</b>	0.00	<b>0.20</b>	0.00	<b>0.40</b>	0.00	0.05	0.00	<b>85869.16</b>	<b>438193.56</b>	<b>387.42</b>	<b>376.47</b>	<b>161.65</b>	<b>189.00</b>	<b>8773.00</b>	<b>5574.45</b>
MLP (Policy Gradient)	<b>0.75</b>	0.00	—	—	—	—	—	—	5.17	400.00	—	—	—	—	—	—
gpt-oss-120b	<b>0.20</b>	0.00	—	—	—	—	—	—	<b>115585.30</b>	<b>197923.80</b>	—	—	—	—	—	—
Grok-4	<b>0.20</b>	0.00	<b>0.05</b>	0.00	0.00	0.00	0.00	0.00	3.75	9.45	3.35	<b>60.00</b>	3.45	3.25	42.00	9.75
Kimi K2 (2024-09-05)	0.15	0.00	—	—	—	—	—	—	47446.30	1297.06	—	—	—	—	—	—
Gemini 2.5 Pro	0.10	0.00	<b>0.05</b>	0.00	<b>0.05</b>	0.00	0.00	0.00	180.50	114.45	36.15	36.25	7.60	8.25	27.35	25.65
DeepSeek-R1 (2025-05-28)	0.05	0.00	—	—	—	—	—	—	28143.85	37819.20	—	—	—	—	—	—
Claude Sonnet 4	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	69.10	126.75	27.55	23.20	6.70	25.10	19.30	65.35
Qwen3-Max	0.05	0.00	—	—	—	—	—	—	35.20	43.35	—	—	—	—	—	—
DeepSeek-V3.1	0.05	0.00	—	—	—	—	—	—	33.80	20.85	—	—	—	—	—	—
doubaod-seed-1-6-vision	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	10.70	11.25	12.75	12.35	12.75	10.65	14.65	13.85
InternVL-3 (78B)	0.00	0.00	0.00	0.00	<b>0.05</b>	0.00	0.00	0.00	<b>56499.75</b>	61.15	<b>48.45</b>	49.75	42.90	38.11	62.00	79.95
Qwen2.5-VL-72B-Instruct	0.00	0.00	0.00	0.00	<b>0.05</b>	0.00	0.00	0.00	47390.10	<b>51351.79</b>	45.25	36.10	13.30	15.55	<b>80.75</b>	30.05
kimi-vl-a3b-thinking	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	5276.13	0.18	0.00	0.11	0.00	0.00	1.06	0.05
GPT-4o	0.00	0.00	0.00	<b>0.10</b>	0.00	0.00	0.00	0.00	83.90	104.10	<b>106.70</b>	<b>104.50</b>	<b>50.40</b>	<b>59.00</b>	<b>118.45</b>	<b>115.35</b>
GLM-4.5V	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	36.85	55.55	30.75	39.70	<b>51.70</b>	<b>58.21</b>	35.50	<b>101.70</b>
Gemma-3-27B-IT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	22.60	18.25	30.30	26.10	19.95	20.30	23.60	16.10
Seed-OSS-36B-Instruct	0.00	0.00	—	—	—	—	—	—	14.68	10.94	—	—	—	—	—	—

#### 4.1 EXPERIMENT 1: BASIC AGENT WITH NO AID

**Experimental Setup.** To answer our first research question, we establish the baseline capabilities of unaided LLMs. For this experiment, we utilize the **Basic Agent** configuration. This agent is provided with only the fundamental interaction tools: `make_move`, `get_observation`, and, for Tier 3 tasks, `apply_view_transformation`. It must solve the task from first principles, relying entirely on its internal faculties for planning and reasoning. We evaluated this agent across all four observation modalities (Full Symbolic, Full Visual, Face view, and Partial Visual) on both short- and long-horizon tasks. The results are presented in Table 5.

**Key Observations.** (1) All models exhibit a 0.00 pass rate on long-horizon tasks across all input modalities. (2) For short-horizon tasks, non-zero pass rates are achieved almost exclusively with the symbolic string input; performance on all visual inputs is near or at zero for most models. (3) A clear performance hierarchy is evident, with GPT-5’s 0.75 pass rate on the symbolic task significantly exceeding all other models. A Policy Gradient agent matches GPT-5’s performance, outperforming most LLMs in this setting. (4) On the Full Symbolic task, a subset of models engage in computationally intensive search, indicated by average #MM counts several orders of magnitude higher than other models.

**Insights.** (1) The universal failure on long-horizon tasks is direct evidence of a fundamental deficit in **Long-Horizon State Tracking through Mental Simulation**. A related case study is presented in Sec. L.2. (2) The sharp performance decline from symbolic to visual inputs indicates that **Visual Thinking** is a primary limiting factor for these agents. (3) While symbolic inputs enable search-based strategies, they are often computationally expensive. A notable phenomenon emerges in these tasks: agents exhibit a diversity of problem-solving strategies. Lower #MM values typically correspond to directly reasoning through the sequence of moves logically, whereas higher #MM values are indicative of search-based strategies. The choice of searching algorithm substantially impacts pass rates. More capable agents, such as GPT-5, tend to systematically search using algorithms like beam search and iterative deepening depth-first search (IDDFS) with skills like backtracking, as shown in Sec. L.1 and L.2. In contrast, less capable agents often devolve into largely unguided enumeration (shown in L.3). However, even models such as GPT-5 struggle to perform effective pruning; their capabilities remain insufficient to curb the rapid growth in computational complexity, leading to failures on long-horizon tasks (shown in L.4).

#### 4.2 EXPERIMENT 2: CUBEAGENT WITH REWARDS

**Experimental Setup.** Our second experiment was designed to measure the impact of different dense reward mechanisms on agent performance. We used the **Basic Agent** agent configuration

378 Table 3: Pass rates of different agent types across modalities and horizons on CubeBench. Metrics  
 379 include **Pass rate** (higher is better). Modalities: Full Symbolic, Full Visual, Face  
 380 View, Vertex View. Tier 3 denotes the hardest split and is evaluated under two projections:  
 381 Face View and Vertex View. Horizons: Short (S) and Long (L).

Model	Reward Type	Full Symbolic		Full Visual		Face view		Vertex view	
		S	L	S	L	S	L	S	L
GPT-5	no reward	0.75	0.00	0.20	0.00	0.40	0.00	0.05	0.00
	face	0.85	0.00	0.55	0.00	0.50	0.00	0.40	0.00
	sticker	0.65	0.00	0.55	0.00	0.55	0.00	0.50	0.00
	heuristic	0.50	0.00	0.45	0.00	0.65	0.00	0.30	0.00
Gemini 2.5 Pro	no reward	0.10	0.00	0.05	0.00	0.05	0.00	0.00	0.00
	face	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	sticker	0.10	0.00	0.00	0.00	0.05	0.00	0.00	0.00
	heuristic	0.05	0.00	0.00	0.00	0.10	0.00	0.00	0.00
Claude Sonnet 4	no reward	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	face	0.10	0.00	0.10	0.00	0.05	0.00	0.00	0.00
	sticker	0.25	0.00	0.15	0.00	0.00	0.00	0.05	0.00
	heuristic	0.20	0.00	0.05	0.00	0.05	0.00	0.10	0.00

394 as the testbed. Its performance was evaluated under four distinct conditions: a baseline with no  
 395 progressive feedback (no reward), and three conditions providing different dense reward signals  
 396 (face, sticker, and heuristic), which are introduced in Sec. 3.1.4 and used as the return  
 397 value for the `make_move` function. This direct comparison aims to clarify whether, and what kind  
 398 of dense rewardss can aid the agent. The results are presented in Table 3.

400 **Key Observations.** (1) On short-horizon tasks, dense rewards generally lead to an increase in pass  
 401 rates. (2) The pass rate on all long-horizon tasks remains at 0.00, regardless of the presence or type  
 402 of dense reward. (3) The impact of rewards is inconsistent; in some cases, such as for GPT-5 on the  
 403 Full Symbolic task with heuristic or sticker rewards, performance is lower than the  
 404 no-reward baseline. (4) The ability to leverage rewards varies notably across models.

405 **Insights.** (1) Dense rewards can guide an agent’s search on short-horizon tasks by providing a local  
 406 heuristic guide. (2) The failure of rewards on long-horizon tasks indicates that local feedback cannot  
 407 compensate for a fundamental deficit in long-horizon state tracking. (3) On visual inputs, agents may  
 408 leverage reward signals through symbolic reasoning, bypassing genuine visual reasoning, as shown  
 409 in Sec. L.5. (4) For more capable agents like GPT-5, an external reward can potentially conflict with  
 410 their emergent internal strategies, leading to suboptimal performance. For less capable agents that  
 411 may lack a strong internal strategy, any form of guidance from a dense reward is often helpful, as  
 412 seen with Claude Sonnet 4. The case studies are presented in Sec. L.6 and L.7 respectively.

### 413 4.3 EXPERIMENT 3: CUBEAGENT WITH SOLVER TOOLS

415 **Experimental Setup.** To precisely identify the primary bottlenecks in agent performance, our final  
 416 experiment removes the burden of long-horizon planning by equipping agents with an optimal  
 417 solver. We introduce two distinct configurations to isolate different cognitive challenges: the  
 418 **Standard-Solver Agent** and the **Ideal-Solver Agent**. The **Standard-Solver Agent** is given a solver  
 419 that requires a specific, strict symbolic input format. To succeed, this agent must first accurately per-  
 420 ceive the cube’s state, then perform the crucial step of **translating** that perception into the required  
 421 format. *This translation process is non-trivial, as it requires spatial understanding to reconcile*  
 422 *potential differences between the environment’s state representation (e.g., one type of 2D unfolded*  
 423 *map) and the solver’s expected input (e.g., a different face order or vertex numbering scheme).* Fi-  
 424 *nally, the agent must execute the solver’s plan. This setup tests the agent’s ability to handle state*  
 425 *reconstruction, spatial transformation, and procedural tool use.*

426 In contrast, the **Ideal-Solver Agent** is provided with a more advanced tool that automates the trans-  
 427 lation step. This agent can directly pass its perceived state to the solver, thus bypassing the data  
 428 formatting challenge. By comparing the performance of these two agents, we can isolate whether  
 429 failures stem from reconstructing a state from perception or from the challenge of spatial under-  
 430 standing during translation. The results are presented in Table 4.

431 **The Diagnosing Framework.** Our evaluation is structured as a progressive, three-step diagnostic  
 432 process designed to systematically isolate and assess the core cognitive faculties of an agent. (1)

432 Table 4: Comparison of pass rates for *Basic*, *Standard-Solver*, and *Ideal-Solver* agent configura-  
 433 tions. Modalities: Full Symbolic, Full Visual, Face View, Vertex View. Tier 3  
 434 denotes the hardest split and is evaluated under two projections: Face View and Vertex View.  
 435 Horizons: Short (S) and Long (L).

Model	Agent Type	Full Symbolic		Full Visual		Face view		Vertex view	
		S	L	S	L	S	L	S	L
GPT-5	Basic	0.75	0.00	0.20	0.00	0.40	0.00	0.05	0.00
	Standard-Solver	0.95	0.95	0.65	0.70	1.00	0.95	0.00	0.00
	Ideal-Solver	1.00	1.00	0.95	0.80	0.85	1.00	0.00	0.00
Gemini 2.5 Pro	Basic	0.10	0.00	0.05	0.00	0.05	0.00	0.00	0.00
	Standard-Solver	0.70	0.65	0.25	0.00	0.20	0.00	0.00	0.00
	Ideal-Solver	1.00	1.00	0.25	0.00	0.00	0.00	0.00	0.00
Claude Sonnet 4	Basic	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	Standard-Solver	0.35	0.85	0.00	0.00	0.00	0.00	0.00	0.00
	Ideal-Solver	1.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00

446  
 447 **Diagnosing Long-Horizon State Tracking and Planning.** We first diagnose long-horizon tracking  
 448 by comparing the *Basic Agent* with the *Standard-Solver Agent* on long-horizon symbolic tasks. The  
 449 *Basic Agent* relies on internal reasoning, while the *Standard-Solver* outsources the planning chal-  
 450 lenge to an optimal tool. Their performance gap reveals the agent’s intrinsic planning capability. (2)  
 451 **Diagnosing Spatial Reasoning and Procedural Tool Use.** Next, we diagnose the spatial reasoning  
 452 required for tool use by comparing the *Standard-Solver Agent* to the *Ideal-Solver Agent*. Since  
 453 both agents offload planning, the performance gap isolates the challenge of spatial thinking, which  
 454 translates perceptual input into a usable format for the tool. (3) **Diagnosing Active Exploration**  
 455 **under Partial Observation.** Finally, to isolate exploration, we evaluate the *Ideal-Solver Agent* in a  
 456 partial observation setting. The ideal tool removes both planning and translation challenges, leaving  
 457 only the task of reconstructing a complete world model from fragmented information. Success here  
 458 depends entirely on the agent’s ability to actively explore its environment.

459 **Key Observations.** (1) The addition of the tools generally leads to marked performance gains  
 460 compared to the basic agent. (2) There is still a performance gap between the Standard-Solver and  
 461 Ideal-Solver agents. (3) On Full Visual and Face view tasks, only GPT-5 maintains strong  
 462 performance, while all models fail universally on the Vertex view task.

463 **Insights.** (1) High-level, multi-step planning, or **Long-Horizon State Tracking**, is a primary deficit  
 464 that can be successfully offloaded to external solvers. (2) The procedural challenge of using tools  
 465 is non-trivial, rendering **Spatial Reasoning** an important challenge to resolve for further develop-  
 466 ment. (3) An unanticipated but noteworthy finding is the emergence of tool-learning strategies in  
 467 the *Standard-Solver Agent*. In some instances, we observed a remarkable capability for autonomous  
 468 tool-learning, where agents learn to master the tool through trial-and-error experimentation for this  
 469 spatial conversation, as shown in Sec. L.8. (4) A significant performance gap exists between the  
 470 Face view and Vertex view tasks. The reason is that the orderly, grid-like structure of the  
 471 Face view allows agents to succeed by recasting the task as an algorithmic parsing problem.  
 472 This indicates that models will attempt to bypass direct spatial reasoning in favor of a parsing-based  
 473 approach whenever possible, and their performance suffers when the input’s complexity, as in the  
 474 Vertex view, makes this bypass strategy infeasible. The corresponding case study is presented  
 475 in Sec. L.9.

## 476 5 CONCLUSION

477 In this work, we introduced CubeBench, a diagnostic benchmark designed to probe the cognitive  
 478 faculties required for **spatial reasoning**. Our comprehensive experiments demonstrate a critical fail-  
 479 ure in current leading models, which uniformly achieve a zero pass rate on all long-horizon tasks  
 480 and struggle to bridge the gap from visual perception to symbolic understanding. Our diagnostic  
 481 framework successfully isolated these bottlenecks, confirming fundamental deficits in Spatial Rea-  
 482 soning, Long-Horizon State Tracking through Mental Simulation, and Exploration and Reasoning  
 483 under Partial Observation. Our findings underscore the need for future research to focus on develop-  
 484 ing more robust spatial mental models and grounding agents in the principles of three-dimensional  
 485 interaction to unlock their potential in the physical world.

486 REPRODUCIBILITY STATEMENT  
487488 We provide necessary tool suite to reproduce our results at <https://anonymous.4open.science/r/CubeBench-ICLR26-34B7/>.  
489490  
491 REFERENCES  
492493 Optimal solutions for the rubik’s cube — kociemba’s algorithm. [https://en.wikipedia.org/wiki/Optimal\\_solutions\\_for\\_the\\_Rubik%27s\\_Cube#Kociemba%27s\\_algorithm](https://en.wikipedia.org/wiki/Optimal_solutions_for_the_Rubik%27s_Cube#Kociemba%27s_algorithm). Accessed September 2025.  
494495  
496 Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel  
497 Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language  
498 model for few-shot learning. *Advances in neural information processing systems*, 35:23716–  
499 23736, 2022.  
500501 Reem Aleithan, Haoran Xue, Mohammad Mahdi Mohajer, Elijah Nnorom, Gias Uddin, and Song  
502 Wang. Swe-bench+: Enhanced coding benchmark for llms. *arXiv preprint arXiv:2410.06992*,  
503 2024.  
504505 Josep Aulinas, Yvan Petillot, Joaquim Salvi, and Xavier Lladó. The slam problem: a survey. *Artificial  
Intelligence Research and Development*, pp. 363–371, 2008.  
506507 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,  
508 Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*,  
509 2025.  
510511 Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio  
512 Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, et al. Mle-bench: Evaluating machine learning  
513 agents on machine learning engineering. *arXiv preprint arXiv:2410.07095*, 2024.  
514515 Kai Chen, Xinfeng Li, Tianpei Yang, Hewei Wang, Wei Dong, and Yang Gao. Mdteamgpt: A  
516 self-evolving llm-based multi-agent framework for multi-disciplinary team medical consultation.  
517 *arXiv preprint arXiv:2503.13856*, 2025a.  
518519 Kaiyuan Chen, Yixin Ren, Yang Liu, Xiaobo Hu, Haotong Tian, Tianbao Xie, Fangfu Liu, Haoye  
520 Zhang, Hongzhang Liu, Yuan Gong, et al. Xbench: Tracking agents productivity scaling with  
521 profession-aligned real-world evaluations. *arXiv preprint arXiv:2506.13651*, 2025b.  
522523 Shiqi Chen, Tongyao Zhu, Ruochen Zhou, Jinghan Zhang, Siyang Gao, Juan Carlos Niebles, Mor  
524 Geva, Junxian He, Jiajun Wu, and Manling Li. Why is spatial reasoning hard for vlms? an  
525 attention mechanism perspective on focus areas. *arXiv preprint arXiv:2503.01773*, 2025c.  
526527 Zhe Chen, Jiannan Wu, Wenhui Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong  
528 Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning  
529 for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer  
530 vision and pattern recognition*, pp. 24185–24198, 2024a.  
531532 Ziru Chen, Shijie Chen, Yuting Ning, Qianheng Zhang, Boshi Wang, Botao Yu, Yifei Li, Zeyi Liao,  
533 Chen Wei, Zitong Lu, et al. Scienceagentbench: Toward rigorous assessment of language agents  
534 for data-driven scientific discovery. *arXiv preprint arXiv:2410.05080*, 2024b.  
535536 Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. Mem0: Building  
537 production-ready ai agents with scalable long-term memory. *arXiv preprint arXiv:2504.19413*,  
538 2025.  
539540 Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su.  
541 Mind2web: Towards a generalist agent for the web. *Advances in Neural Information Processing  
542 Systems*, 36:28091–28114, 2023.  
543544 Jinyuan Fang, Yanwen Peng, Xi Zhang, Yingxu Wang, Xinhao Yi, Guibin Zhang, Yi Xu, Bin Wu,  
545 Siwei Liu, Zihao Li, et al. A comprehensive survey of self-evolving ai agents: A new paradigm  
546 bridging foundation models and lifelong agentic systems. *arXiv preprint arXiv:2508.07407*, 2025.  
547

540     ARC Prize Foundation. Arc-agi-3: Interactive reasoning benchmark, 2025. URL <https://arcprize.org/arc-agi/3/>. Accessed: 2025-09-20.  
541  
542

543     Chen Gao, Baining Zhao, Weichen Zhang, Jinzhu Mao, Jun Zhang, Zhiheng Zheng, Fanhang Man,  
544     Jianjie Fang, Zile Zhou, Jinqiang Cui, et al. Embodiedcity: A benchmark platform for embodied  
545     agent in real-world city environment. *arXiv preprint arXiv:2410.09604*, 2024.

546     Huan-ang Gao, Jiayi Geng, Wenyue Hua, Mengkang Hu, Xinzhe Juan, Hongzhang Liu, Shilong  
547     Liu, Jiahao Qiu, Xuan Qi, Yiran Wu, et al. A survey of self-evolving agents: On path to artificial  
548     super intelligence. *arXiv preprint arXiv:2507.21046*, 2025.

549  
550     Zhenyu Guan, Xiangyu Kong, Fangwei Zhong, and Yizhou Wang. Richelieu: Self-evolving llm-  
551     based agents for ai diplomacy. *Advances in Neural Information Processing Systems*, 37:123471–  
552     123497, 2024.

553     Mohd Ariful Haque, Justin Williams, Sunzida Siddique, Md. Hujaifa Islam, Hasmot Ali,  
554     Kishor Datta Gupta, and Roy George. Advanced tool learning and selection system (ATLASS):  
555     A closed-loop framework using LLM. *CoRR*, abs/2503.10071, 2025.

556  
557     Mengkang Hu, Pu Zhao, Can Xu, Qingfeng Sun, Jianguang Lou, Qingwei Lin, Ping Luo, and  
558     Saravan Rajmohan. Agentgen: Enhancing planning abilities for large language model based  
559     agent via environment and task generation. *arXiv preprint arXiv:2408.00764*, 2024a.

560     Mengkang Hu, Yuhang Zhou, Wendong Fan, Yuzhou Nie, Bowei Xia, Tao Sun, Ziyu Ye, Zhaoxuan  
561     Jin, Yingru Li, Qiguang Chen, et al. Owl: Optimized workforce learning for general multi-agent  
562     assistance in real-world task automation. *arXiv preprint arXiv:2505.23885*, 2025.

563     Yue Hu, Yuzhu Cai, Yixin Du, Xinyu Zhu, Xiangrui Liu, Zijie Yu, Yuchen Hou, Shuo Tang, and  
564     Siheng Chen. Self-evolving multi-agent collaboration networks for software development. *arXiv  
565     preprint arXiv:2410.16946*, 2024b.

566  
567     Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-  
568     trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint  
569     arXiv:2410.21276*, 2024.

570     Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik  
571     Narasimhan. Swe-bench: Can language models resolve real-world github issues? *arXiv preprint  
572     arXiv:2310.06770*, 2023.

573     Philip N Johnson-Laird. Mental models in cognitive science. *Cognitive science*, 4(1):71–115, 1980.

574  
575     Philip Nicholas Johnson-Laird. *Mental models: Towards a cognitive science of language, inference,  
576     and consciousness*. Number 6. Harvard University Press, 1983.

577     Herbert Kociemba. Cube explorer and kociemba’s two-phase algorithm. <https://kociemba.org/cube.htm>. Accessed September 2025.

578  
579     Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt  
580     Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, et al. Ai2-thor: An interactive 3d environment  
581     for visual ai. *arXiv preprint arXiv:1712.05474*, 2017.

582  
583     Richard E. Korf. Depth-first iterative-deepening: An optimal admissible tree search. *Artificial  
584     Intelligence*, 27(1):97–109, 1985.

585  
586     Richard E. Korf. Finding optimal solutions to rubik’s cube using pattern databases. In *AAAI*, 1997.

587  
588     Ido Levy, Ben Wiesel, Sami Marreed, Alon Oved, Avi Yaeli, and Segev Shlomov. St-  
589     webagentbench: A benchmark for evaluating safety and trustworthiness in web agents. *arXiv  
590     preprint arXiv:2410.06703*, 2024.

591     Bowen Li, Zhaoyu Li, Qiwei Du, Jinqi Luo, Wenshan Wang, Yaqi Xie, Simon Stepputtis, Chen  
592     Wang, Katia P. Sycara, Pradeep Kumar Ravikumar, Alexander G. Gray, Xujie Si, and Sebastian  
593     Scherer. Logicity: Advancing neuro-symbolic ai with abstract urban simulation, 2025. URL  
<https://arxiv.org/abs/2411.00773>.

594 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image  
 595 pre-training with frozen image encoders and large language models. In *International conference*  
 596 *on machine learning*, pp. 19730–19742. PMLR, 2023.

597

598 Manling Li, Shiyu Zhao, Qineng Wang, Kangrui Wang, Yu Zhou, Sanjana Srivastava, Cem Gokmen,  
 599 Tony Lee, Erran Li Li, Ruohan Zhang, et al. Embodied agent interface: Benchmarking llms for  
 600 embodied decision making. *Advances in Neural Information Processing Systems*, 37:100428–  
 601 100534, 2024.

602 Xuechen Liang, Yangfan He, Yinghui Xia, Xinyuan Song, Jianhui Wang, Meiling Tao, Li Sun,  
 603 Xinhang Yuan, Jiayi Su, Keqin Li, Siyuan Chen, and Tianyu Shi. Self-evolving agents with  
 604 reflective and memory-augmented abilities. *arXiv preprint arXiv:2409.00872*, 2024a.

605

606 Xuechen Liang, Yangfan He, Yinghui Xia, Xinyuan Song, Jianhui Wang, Meiling Tao, Li Sun, Xinhang  
 607 Yuan, Jiayi Su, Keqin Li, et al. Self-evolving agents with reflective and memory-augmented  
 608 abilities. *arXiv preprint arXiv:2409.00872*, 2024b.

609

610 Fangzheng Liu, Don Derek Haddad, and Joe Paradiso. Mindcube: an interactive device for gauging  
 611 emotions. In *Adjunct Proceedings of the 37th Annual ACM Symposium on User Interface Software*  
 612 *and Technology*, pp. 1–2, 2024.

613

614 Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding,  
 615 Kaiwen Men, Kejuna Yang, et al. Agentbench: Evaluating llms as agents. *arXiv preprint*  
 616 *arXiv:2308.03688*, 2023.

617 Junyu Luo, Weizhi Zhang, Ye Yuan, Yusheng Zhao, Junwei Yang, Yiyang Gu, Bohan Wu, Binqi  
 618 Chen, Ziyue Qiao, Qingqing Long, Rongcheng Tu, Xiao Luo, Wei Ju, Zhiping Xiao, Yifan Wang,  
 619 Meng Xiao, Chenwu Liu, Jingyang Yuan, Shichang Zhang, Yiqiao Jin, Fan Zhang, Xian Wu,  
 620 Hanqing Zhao, Dacheng Tao, Philip S. Yu, and Ming Zhang. Large language model agent: A sur-  
 621 vey on methodology, applications and challenges, 2025. URL <https://arxiv.org/abs/2503.21460>.

622

623 Xiaowen Ma, Chenyang Lin, Yao Zhang, Volker Tresp, and Yunpu Ma. Agentic neural networks:  
 624 Self-evolving multi-agent systems via textual backpropagation. *arXiv preprint arXiv:2506.09046*,  
 625 2025.

626

627 Grégoire Mialon, Clémentine Fourrier, Thomas Wolf, Yann LeCun, and Thomas Scialom. Gaia:  
 628 a benchmark for general ai assistants. In *The Twelfth International Conference on Learning*  
 629 *Representations*, 2023.

630

631 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and  
 632 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications*  
 633 *of the ACM*, 65(1):99–106, 2021.

634

635 Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, Hugh Zhang, Chen Bo Calvin  
 636 Zhang, Mohamed Shaaban, John Ling, Sean Shi, et al. Humanity’s last exam. *arXiv preprint*  
 637 *arXiv:2501.14249*, 2025.

638

639 Renjie Pi, Jianshu Zhang, Jipeng Zhang, Rui Pan, Zhekai Chen, and Tong Zhang. Image textual-  
 640 ization: An automatic framework for creating accurate and detailed image descriptions. *arXiv*  
 641 *preprint arXiv:2406.07502*, 2024.

642

643 Xavi Puig, Eric Undersander, Andrew Szot, Mikael Dallaire Cote, Ruslan Partsey, Jimmy Yang,  
 644 Ruta Desai, Alexander William Clegg, Michal Hlavac, Tiffany Min, Theo Gervet, Vladimir Von-  
 645 drus, Vincent-Pierre Berges, John Turner, Oleksandr Maksymets, Zsolt Kira, Mrinal Kalakri-  
 646 shnan, Jitendra Malik, Devendra Singh Chaplot, Unnat Jain, Dhruv Batra, Akshara Rai, and  
 647 Roozbeh Mottaghi. Habitat 3.0: A co-habitat for humans, avatars and robots, 2023.

648

649 Jianing Qi, Jiawei Liu, Hao Tang, and Zhigang Zhu. Beyond semantics: Rediscovering spatial  
 650 awareness in vision-language models. *arXiv preprint arXiv:2503.17349*, 2025.

648 Jiahao Qiu, Xuan Qi, Tongcheng Zhang, Xinzhe Juan, Jiacheng Guo, Yifu Lu, Yimin Wang, Zixin  
 649 Yao, Qihan Ren, Xun Jiang, Xing Zhou, Dongrui Liu, Ling Yang, Yue Wu, Kaixuan Huang,  
 650 Shilong Liu, Hongru Wang, and Mengdi Wang. Alita: Generalist agent enabling scalable agentic  
 651 reasoning with minimal predefinition and maximal self-evolution. *CoRR*, abs/2505.20286, 2025.  
 652

653 Chang Qu, Sunhao Dai, Xiaochi Wei, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, Jun Xu, and  
 654 Ji-Rong Wen. From exploration to mastery: Enabling llms to master tools via self-driven interac-  
 655 tions. In *ICLR*. OpenReview.net, 2025.

656 Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel  
 657 Todorov, and Sergey Levine. Learning complex dexterous manipulation with deep reinforcement  
 658 learning and demonstrations. *arXiv preprint arXiv:1709.10087*, 2017.

659 Tomas Rokicki, Herbert Kociemba, Morley Davidson, and John Dethridge. God’s Number is 20.  
 660 <https://www.cube20.org/>, 2010.

661 Tomas Rokicki, Herbert Kociemba, Morley Davidson, and John Dethridge. The diameter of the  
 662 rubik’s cube group is twenty. *SIAM Review*, 56(4):645–670, 2014. doi: 10.1137/140973499.

663 Ravi Sapkota, Konstantinos I Roumeliotis, Shiva Pokhrel, and Manoj Karkee. From self-learning to  
 664 self-evolving architectures in large language models: A short survey. *Authorea Preprints*, 2025.

665 Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain,  
 666 Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A  
 667 Platform for Embodied AI Research. In *Proceedings of the IEEE/CVF International Conference*  
 668 *on Computer Vision (ICCV)*, 2019a.

669 Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain,  
 670 Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied  
 671 ai research. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp.  
 672 9339–9347, 2019b.

673 Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi,  
 674 Luke Zettlemoyer, and Dieter Fox. Alfred: A benchmark for interpreting grounded instructions  
 675 for everyday tasks, 2020. URL <https://arxiv.org/abs/1912.01734>.

676 Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Côté, Yonatan Bisk, Adam Trischler, and Matthew  
 677 Hausknecht. Alfworld: Aligning text and embodied environments for interactive learning, 2021.  
 678 URL <https://arxiv.org/abs/2010.03768>.

679 Sanjana Srivastava, Chengshu Li, Michael Lingelbach, Roberto Martín-Martín, Fei Xia, Kent  
 680 Vainio, Zheng Lian, Cem Gokmen, Shyamal Buch, C. Karen Liu, Silvio Savarese, Hyowon  
 681 Gweon, Jiajun Wu, and Li Fei-Fei. Behavior: Benchmark for everyday household activities  
 682 in virtual, interactive, and ecological environments, 2021. URL <https://arxiv.org/abs/2108.03332>.

683 Hongjin Su, Ruoxi Sun, Jinsung Yoon, Pengcheng Yin, Tao Yu, and Sercan Ö. Arik. Learn-by-  
 684 interact: A data-centric framework for self-adaptive agents in realistic environments, 2025. URL  
 685 <https://arxiv.org/abs/2501.10893>.

686 Andrew Szot, Alex Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre,  
 687 Mustafa Mukadam, Devendra Chaplot, Oleksandr Maksymets, Aaron Gokaslan, Vladimir Von-  
 688 drus, Sameer Dharur, Franziska Meier, Wojciech Galuba, Angel Chang, Zsolt Kira, Vladlen  
 689 Koltun, Jitendra Malik, Manolis Savva, and Dhruv Batra. Habitat 2.0: Training home assistants to  
 690 rearrange their habitat. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

691 Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Bud-  
 692 den, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, et al. Deepmind control suite. *arXiv*  
 693 *preprint arXiv:1801.00690*, 2018.

694 Karthik Valmecikam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kamb-  
 695 hampati. Planbench: An extensible benchmark for evaluating large language models on planning  
 696 and reasoning about change. *Advances in Neural Information Processing Systems*, 36:38975–  
 697 38987, 2023.

702 Hongru Wang, Yujia Qin, Yankai Lin, Jeff Z. Pan, and Kam-Fai Wong. Empowering large lan-  
 703 guage models: Tool learning for real-world interaction. In *Proceedings of the 47th Interna-*  
 704 *tional ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR*  
 705 *'24*, pp. 2983–2986, New York, NY, USA, 2024a. Association for Computing Machinery. ISBN  
 706 9798400704314. doi: 10.1145/3626772.3661381. URL <https://doi.org/10.1145/3626772.3661381>.

708 Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Ji-  
 709 akai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Jirong Wen. A survey on  
 710 large language model based autonomous agents. *Frontiers of Computer Science*, 18(6), March  
 711 2024b. ISSN 2095-2236. doi: 10.1007/s11704-024-40231-1. URL <http://dx.doi.org/10.1007/s11704-024-40231-1>.

713 Renxi Wang, Xudong Han, Lei Ji, Shu Wang, Timothy Baldwin, and Haonan Li. Toolgen: Unified  
 714 tool retrieval and calling via generation. In *ICLR*. OpenReview.net, 2025.

716 Jason Wei, Zhiqing Sun, Spencer Papay, Scott McKinney, Jeffrey Han, Isa Fulford, Hyung Won  
 717 Chung, Alex Tachard Passos, William Fedus, and Amelia Glaese. Browsecmp: A simple yet  
 718 challenging benchmark for browsing agents. *arXiv preprint arXiv:2504.12516*, 2025.

719 Jialong Wu, Wenbiao Yin, Yong Jiang, Zhenglin Wang, Zekun Xi, Runnan Fang, Linhai Zhang,  
 720 Yulan He, Deyu Zhou, Pengjun Xie, et al. Webwalker: Benchmarking llms in web traversal.  
 721 *arXiv preprint arXiv:2501.07572*, 2025.

723 Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh J Hua,  
 724 Zhoujun Cheng, Dongchan Shin, Fangyu Lei, et al. Osworld: Benchmarking multimodal agents  
 725 for open-ended tasks in real computer environments. *Advances in Neural Information Processing  
 726 Systems*, 37:52040–52094, 2024.

727 Frank F Xu, Yufan Song, Boxuan Li, Yuxuan Tang, Kritanjali Jain, Mengxue Bao, Zora Z Wang,  
 728 Xuhui Zhou, Zhitong Guo, Murong Cao, et al. Theagentcompany: benchmarking llm agents on  
 729 consequential real world tasks. *arXiv preprint arXiv:2412.14161*, 2024.

730 Wentui Xu, Dalin Lyu, Weihang Wang, Jie Feng, Chen Gao, and Yong Li. Defining and evaluating  
 731 visual language models' basic spatial abilities: A perspective from psychometrics. *arXiv preprint  
 732 arXiv:2502.11859*, 2025a.

734 Wujiang Xu, Kai Mei, Hang Gao, Juntao Tan, Zujie Liang, and Yongfeng Zhang. A-mem: Agentic  
 735 memory for llm agents. *arXiv preprint arXiv:2502.12110*, 2025b.

736 Jihan Yang, Shusheng Yang, Anjali W Gupta, Rilyn Han, Li Fei-Fei, and Saining Xie. Thinking in  
 737 space: How multimodal large language models see, remember, and recall spaces. In *Proceedings  
 738 of the Computer Vision and Pattern Recognition Conference*, pp. 10632–10643, 2025a.

739 John Yang, Carlos E Jimenez, Alex L Zhang, Kilian Lieret, Joyce Yang, Xindi Wu, Ori Press,  
 740 Niklas Muennighoff, Gabriel Synnaeve, Karthik R Narasimhan, et al. Swe-bench multimodal:  
 741 Do ai systems generalize to visual software domains? *arXiv preprint arXiv:2410.03859*, 2024.

743 Rui Yang, Hanyang Chen, Junyu Zhang, Mark Zhao, Cheng Qian, Kangrui Wang, Qineng Wang,  
 744 Teja Venkat Koripella, Marziyeh Movahedi, Manling Li, et al. Embodiedbench: Compre-  
 745 hensive benchmarking multi-modal large language models for vision-driven embodied agents. *arXiv  
 746 preprint arXiv:2502.09560*, 2025b.

747 Jianzhu Yao, Kevin Wang, Ryan Hsieh, Haisu Zhou, Tianqing Zou, Zerui Cheng, Zhangyang Wang,  
 748 and Pramod Viswanath. Spin-bench: How well do llms plan strategically and reason socially?,  
 749 2025. URL <https://arxiv.org/abs/2503.12349>.

750 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.  
 751 React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*,  
 752 2022.

754 Hongli Yu, Tinghong Chen, Jiangtao Feng, Jiangjie Chen, Weinan Dai, Qiying Yu, Ya-Qin Zhang,  
 755 Wei-Ying Ma, Jingjing Liu, Mingxuan Wang, et al. Memagent: Reshaping long-context llm with  
 multi-conv rl-based memory agent. *arXiv preprint arXiv:2507.02259*, 2025.

756 Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey  
 757 Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning.  
 758 In *Conference on robot learning*, pp. 1094–1100. PMLR, 2020.

759

760 Danyang Zhang, Lu Chen, Situo Zhang, Hongshen Xu, Zihan Zhao, and Kai Yu. Large language  
 761 models are semi-parametric reinforcement learning agents. *Advances in Neural Information Pro-*  
 762 *cessing Systems*, 36, 2024a.

763 Dingchu Zhang, Yida Zhao, Jialong Wu, Baixuan Li, Wenbiao Yin, Liwen Zhang, Yong Jiang,  
 764 Yufeng Li, Kewei Tu, Pengjun Xie, and Fei Huang. Evolvesearch: An iterative self-evolving  
 765 search agent. *arXiv preprint arXiv:2505.22501*, 2025a.

766

767 Genghan Zhang, Weixin Liang, Olivia Hsu, and Kunle Olukotun. Adaptive self-improvement llm  
 768 agentic system for ml library development. *arXiv preprint arXiv:2502.02534*, 2025b.

769

770 Guanghui Zhang, Li Niu, Jianwei Fang, Kun Wang, Lu Bai, et al. Multi-agent architecture search  
 771 via agentic supernet. *arXiv preprint arXiv:2502.04180*, 2025c.

772

773 Wenyu Zhang, Wei En Ng, Lixin Ma, Yuwen Wang, Junqi Zhao, Allison Koenecke, Boyang Li, and  
 774 Lu Wang. Sphere: Unveiling spatial blind spots in vision-language models through hierarchical  
 775 evaluation. *arXiv preprint arXiv:2412.12693*, 2024b.

776

777 Zheyuan Zhang, Fengyuan Hu, Jayjun Lee, Freda Shi, Parisa Kordjamshidi, Joyce Chai, and Ziqiao  
 778 Ma. Do vision-language models represent space and how? evaluating spatial frame of reference  
 779 under ambiguities. *arXiv preprint arXiv:2410.17385*, 2024c.

780

781 Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. Expel: Llm  
 782 agents are experiential learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*,  
 783 volume 38, pp. 19632–19642, 2024a.

784

785 Haiteng Zhao, Chang Ma, Guoyin Wang, Jing Su, Lingpeng Kong, Jingjing Xu, Zhi-Hong Deng,  
 786 and Hongxia Yang. Empowering large language model agents through action learning. *CoRR*,  
 787 abs/2402.15809, 2024b.

788

789 Boyuan Zheng, Michael Y. Fatemi, Xiaolong Jin, Zora Zhiruo Wang, Apurva Gandhi, Yueqi Song,  
 790 Yu Gu, Jayanth Srinivasa, Gaowen Liu, Graham Neubig, and Yu Su. Skillweaver: Web agents  
 791 can self-improve by discovering and honing skills. *CoRR*, abs/2504.07079, 2025.

792

793 Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng,  
 794 Tianyue Ou, Yonatan Bisk, Daniel Fried, et al. Webarena: A realistic web environment for build-  
 795 ing autonomous agents. *arXiv preprint arXiv:2307.13854*, 2023.

796

797 Wangchunshu Zhou, Yixin Ou, Shengwei Ding, Long Li, Jialong Wu, Tiannan Wang, Jiamin Chen,  
 798 Shuai Wang, Xiaohua Xu, Ningyu Zhang, et al. Symbolic learning enables self-evolving agents.  
 799 *arXiv preprint arXiv:2406.18532*, 2024.

800

801 Kunlun Zhu, Hongyi Du, Zhaochen Hong, Xiaocheng Yang, Shuyi Guo, Zhe Wang, Zhenhailong  
 802 Wang, Cheng Qian, Xiangru Tang, Heng Ji, et al. Multiagentbench: Evaluating the collaboration  
 803 and competition of llm agents. *arXiv preprint arXiv:2503.01935*, 2025.

804

805 Yangyang Zhuang, Wenjia Jiang, Jiayu Zhang, Ze Yang, Joey Tianyi Zhou, and Chi Zhang.  
 806 Learning to be a doctor: Searching for effective medical agent architectures. *arXiv preprint*  
 807 *arXiv:2504.11301*, 2025.

808

809

810 A LLM USAGE STATEMENT  
811812 *Model.* The LLMs employed in our study are GPT-5 and Gemini 2.5 Pro.  
813814 *Scope.* We used large language models (LLMs) only as general-purpose assistants for language pol-  
815 ishing, typo checking, and minor code boilerplate generation. LLMs did not contribute to research  
816 ideation or produce novel scientific claims, proofs, or results.817 *Human oversight and verification.* All text and code produced with LLM assistance were reviewed,  
818 corrected, and verified by the authors. Experimental results were reproduced independently of any  
819 LLM outputs.820 *Data governance.* We did not share proprietary or sensitive data with third-party services beyond  
821 materials already included in the anonymous submission artifacts.822 *Attribution.* LLMs are not authors and bear no responsibility for the content; full responsibility lies  
823 with the paper’s authors.  
824825 B WHAT EXACTLY DOES CUBE BENCH MEASURE?  
826827 CubeBench is not intended to be a broad coverage benchmark that competes with large embodied  
828 suites or classical Rubik’s Cube solvers. Instead, it serves as a *minimal, verifiable, factorized* diag-  
829 nóstic environment. By stripping away perception noise, multi-object dynamics, affordances, and  
830 actuation, it allows us to focus on three core cognitive abilities that repeatedly emerge as bottle-  
831 necks for LLM/MLLM agents in more complex settings: (i) 3D spatial reasoning, (ii) long-horizon  
832 non-commutative planning, and (iii) belief-state construction under partial observability.  
833834 Our three-tier design decouples these abilities by progressively increasing the burden placed on  
835 the agent’s internal world model. Tier 1 (*Full Symbolic*) exposes a complete 54-character state, so  
836 all perception and grounding are provided by the environment; this setting primarily stresses non-  
837 commutative long-horizon planning and state tracking. Tier 2 (*Full Visual*) replaces the symbolic  
838 string with a 2D unfolded image, forcing the agent to construct its own symbolic representation  
839 from pixels: segmenting the cube, clustering colors, assigning stickers to faces, and mapping the  
840 2D layout to a consistent 3D frame before any planning can begin. Tier 3 (*Partial Visual*) further  
841 restricts each observation to a single face or corner view, plus view-change actions, so the agent  
842 must actively explore, aggregate partial views over time, and maintain a coherent latent world state  
843 while its own actions continuously perturb the cube.844 On top of these observation tiers, the different solver configurations act as *controls* that selectively  
845 remove specific difficulties. A *Basic Agent* must handle perception, internal state tracking, planning,  
846 and formatting. In the *Standard-Solver* setting, we outsource optimal planning to a Kociembab-  
847 based solver, so performance is driven mainly by visual → symbolic translation and correct schema  
848 formatting. The *Ideal-Solver* setting goes one step further by also hiding the solver’s input schema,  
849 leaving only the requirement to output a correct symbolic cube state; any gap between Standard-  
850 and Ideal-Solver performance therefore isolates visual grounding and spatial mapping errors rather  
851 than planning or string-format issues.852 This factorized design explains “what” CubeBench measures in contrast to existing planning and  
853 embodied benchmarks. Classical planning suites such as PlanBench (Valmeekam et al., 2023),  
854 SPIN-Bench (Yao et al., 2025), and ARC/ARC-AGI (Foundation, 2025) probe rich algorithmic and  
855 combinatorial structures, but they do not directly target the specific triad of 3D spatial reasoning,  
856 non-commutative dynamics, and partial-observation belief-state tracking in a single, fully verifi-  
857 able physical system. Conversely, embodied 3D suites such as ALFRED (Shridhar et al., 2020),  
858 Habitat(Savva et al., 2019a; Szot et al., 2021; Puig et al., 2023), BEHAVIOR-100 (Srivastava et al.,  
859 2021), LogiCity (Li et al., 2025), EmbodiedBench (Yang et al., 2025b), and EAI (Li et al., 2024)  
860 place agents in visually rich, multi-object worlds with realistic physics and affordances, but neces-  
861 sarily entangle perception, control, and high-level reasoning. As several of these works themselves  
862 emphasize, final success rates in such environments make it difficult to pinpoint which cognitive  
863 ability has failed.864 CubeBench occupies an orthogonal niche in this landscape. It uses a single rigid object with de-  
865 terministic kinematics to provide a low-noise, automatically verifiable testbed where failures can

be crisply attributed: incorrect 3D world modeling, short effective planning horizons, or unstable belief-state tracking under partial observation. The negative results we obtain—universal collapse at depth 8, severe degradation from symbols to images, and near-zero success under partial views—mirror the failure modes reported in broader embodied benchmarks, but in a setting where they can be disentangled and systematically ablated. In this sense, CubeBench is best viewed as a *first diagnostic stop* before running expensive embodied evaluations: if an agent already fails to maintain a consistent 3D mental model and long-horizon plan in this simplified domain, it is unlikely to succeed in more complex physical worlds.

### C REVISED BASELINE PERFORMANCE USING MOVE RATIOS

To account for the variation in task difficulty, we introduce the Number of Move Ratio (#MR) as a normalized measure of search efficiency. Optimal path lengths increase non-linearly with scramble depth, therefore we define #MR as the ratio of the agent’s move count to the optimal solution length:

$$\text{#MR (number of move ratio)} = \frac{\text{#MM (number of make\_moves)}}{\text{depth (number of optimal moves)}}$$

Table 5: Baseline performance across modalities and horizons on CubeBench. Top row: metric groups (**Pass rate**, **#MR**, where #MR is the average number of move ratios; second row: observation modalities; third row: task horizons (Short = S, depths 1–4; Long = L, depths 8, 12, 16, 20). **Tier 3** denotes the hardest split, evaluated under two projections: *Face View* and *Vertex View*. Blue shading denotes open-source models, and pink denotes proprietary models. For each metric column, we shade the top-3 entries (red = 1st, orange = 2nd, yellow = 3rd). We also train an MLP with policy gradient on the Full Symbolic setting; details are in Sec. H. “—” : Model does not support visual inputs.

Model	Pass rate								#MR							
	Full Symbolic		Full Visual		Face view		Vertex view		Full Symbolic		Full Visual		Face view		Vertex view	
	S	L	S	L	S	L	S	L	S	L	S	L	S	L	S	L
GPT-5	0.75	0.00	0.20	0.00	0.40	0.00	0.05	0.00	27124.66	33982.78	149.29	29.22	51.25	17.69	8378.55	672.57
MLP (Policy Gradient)	0.75	0.00	—	—	—	—	—	—	5.17	400.00	—	—	—	—	—	—
gpt-oss-120b	0.20	0.00	—	—	—	—	—	—	47173.21	41219.57	—	—	—	—	—	—
Grok-4	0.20	0.00	0.05	0.00	0.00	0.00	0.00	0.00	1.84	0.97	1.76	5.28	1.81	0.25	18.36	0.84
Kimi K2 (2024-09-05)	0.15	0.00	—	—	—	—	—	—	23705.38	118.25	—	—	—	—	—	—
Gemini 2.5 Pro	0.10	0.00	0.05	0.00	0.05	0.00	0.00	0.00	59.01	12.55	19.30	3.13	4.05	0.55	13.28	1.85
DeepSeek-R1 (2025-05-28)	0.05	0.00	—	—	—	—	—	—	14074.89	4724.51	—	—	—	—	—	—
Claude Sonnet 4	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	28.59	9.05	14.81	1.93	2.67	2.43	9.40	5.75
Qwen3-Max	0.05	0.00	—	—	—	—	—	—	15.92	3.48	—	—	—	—	—	—
DeepSeek-V3.1	0.05	0.00	—	—	—	—	—	—	21.89	1.66	—	—	—	—	—	—
doubao-seed-1-6-vision	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	4.77	0.94	6.79	1.00	6.26	0.84	7.47	1.10
InternVL-3 (78B)	0.00	0.00	0.00	0.00	0.05	0.00	0.00	0.00	56466.74	4.85	23.27	3.95	18.16	2.61	29.14	6.84
Qwen2.5-VL-72B-Instruct	0.00	0.00	0.00	0.00	0.05	0.00	0.00	0.00	43712.71	2575.65	18.26	2.54	5.97	1.26	27.28	2.44
kimi-vl-a3b-thinking	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2623.57	0.01	0.00	0.01	0.00	0.00	0.36	0.00
GPT-4o	0.00	0.00	0.00	0.00	0.10	0.00	0.00	0.00	48.10	8.01	52.67	8.48	24.85	4.55	55.41	9.41
GLM-4.5V	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	16.61	5.44	15.48	3.35	28.49	5.17	21.08	8.33
Gemma-3-27B-IT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	12.29	1.56	15.49	2.09	9.38	1.64	11.49	1.21
Seed-OSS-36B-Instruct	0.00	0.00	—	—	—	—	—	—	9.22	0.93	—	—	—	—	—	—

### D THE NO-CODE EXPERIMENT

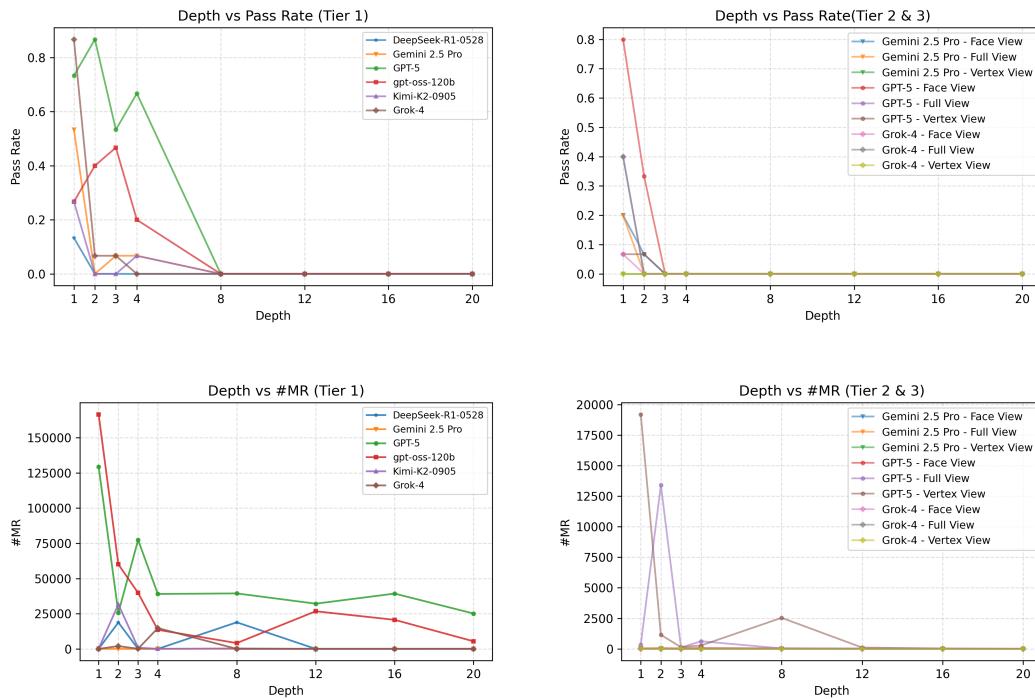
**Experimental Setup.** In our primary evaluation (the **Code** setting, which corresponds to the **Basic Agent** configuration in Experiment 1), agents are permitted to write and execute Python code. While writing code to perform search algorithms is a valid problem-solving strategy in the agentic era, we acknowledge that distinguishing intrinsic spatial reasoning from programmatic search is crucial. To isolate the model’s internal state tracking and planning capabilities from code-based search, we introduce a **No-Code** evaluation mode. In this mode, the agent must output a pre-planned sequence (e.g., “moves”: [‘D’, ‘B’, ‘B’]) without the ability to manage control flows during tool use. This mode is more consistent with common tool-use interfaces such as the OpenAI API. We conducted an ablation study comparing the **No-Code** and **Code** settings on the Short-Horizon Full Symbolic tasks. The results are presented in Table 6.

**Key Observations & Insights.** (1) Models employing intensive search strategies (High #MR) suffered catastrophic declines. Both GPT-5 (0.75 to 0.25) and gpt-oss-120b (0.20 to 0.00) saw their performance evaporate as their search volume collapsed (e.g., GPT-5 #MR: ~27k to ~72; gpt-oss-120b #MR: ~47k to ~14). This confirms that for these agents, the code-based search compensates

918 for limited internal planning. (2) Models with initially low search volume (Low #MR) displayed  
 919 mixed outcomes rather than a uniform drop. Grok-4 improved (0.20 to 0.30) in the No-Code setting,  
 920 while many remained poor.

922 Table 6: Comparison of agent performance in **No-Code** vs. **Code** settings on Short-Horizon Full  
 923 Symbolic tasks. #MR is averaged over testcases.

Model	No Code		Code	
	Pass Rate	#MR	Pass Rate	#MR
GPT-5	0.25	72.79	<b>0.75</b>	27124.66
gpt-oss-120b	0.00	14.20	0.20	47173.21
Grok-4	<b>0.30</b>	2.02	0.20	1.84
Kimi-K2-0905	0.00	55.15	0.15	23705.38
Gemini 2.5 Pro	0.15	18.68	0.10	59.01
Claude Sonnet 4	0.00	68.03	0.05	28.59
DeepSeek-V3.1	0.00	17.11	0.05	21.89
DeepSeek-R1-0528	0.00	12.28	0.05	14074.89
Qwen3-Max	0.00	18.11	0.05	15.92
GPT-4o	0.05	61.53	0.00	48.10
GLM-4.5v	0.00	4.31	0.00	16.61



961 Figure 6: **Pass rates and average number of move ratios (#MR) across varying depths.** Top row:  
 962 Pass Rates for Tier 1 (Left) and Tier 2/3 (Right). Bottom row: Corresponding #MR analysis.

## E DEPTH-WISE METRICS

967 **Experimental Setting.** We adopt the experimental protocol from Experiment 1, extending the eval-  
 968 uation to three independent trials per test case to ensure robustness. For this analysis, we mainly  
 969 focus exclusively on the top-performing models in Experiment 1. The comparative results are pre-  
 970 sented in Figure 6.

971 **Key Observations & Insights.**

(1) *Depth-vs-Performance*. In the Tier 1 setting, top-tier models capable of systematic search (e.g., GPT-5) maintain high accuracy across Depths 1–4, with pass rates decaying gracefully. Conversely, less capable models (e.g., Grok-4) succeed at Depth 1 via direct policy but degrade immediately at Depth 2. In Tier 2/3 settings, the performance cliff arrives earlier. While agents can manage Depth 1 tasks, we observe a drop starting at just Depth 2. A universal limit appears between Depth 4 and 8. Pass rates drop to 0.00 at Depth 8 across all settings, confirming that current planning capabilities fail to track state over extended horizons.

(2) *Depth-vs-#MR*. We interpret the #MR metric not as a measure of “effort” scaling linearly with difficulty, but as an indicator of cognitive mode. A low move count typically signals reliance on direct intuition or trivial heuristics, while a high move count signals the activation of explicit search. The switch is often a discrete jump rather than a gradual increase. A close examination of execution traces reveals how agents dynamically adapt within a single trial. Typically, search-based agents begin with low-cost heuristics, switching to high-volume search when those initial attempts fail. In rare instances (e.g., Appendix L.2), we observe a further strategic shift: abandoning search entirely to implement structured human algorithms, such as the Beginner’s Method (see Appendix J), as a fallback.

## F RELATED WORKS (CONT.)

**Spatial Cognition.** The ability to reason about three-dimensional space (Li et al., 2023; Pi et al., 2024; Alayrac et al., 2022), a cornerstone of intelligence, relies on an internal *spatial mental model* (Johnson-Laird, 1980; 1983) to infer unseen properties and predict the consequences of actions. (Yang et al., 2025a; Zhang et al., 2024c) This concept is rooted in cognitive science and has been a long-standing goal, with specialized systems like SLAM in robotics (Aulinas et al., 2008) and NeRFs (Mildenhall et al., 2021) in computer vision designed to construct explicit 3D representations. However, the intrinsic ability of Large Language Models (LLMs) (Bai et al., 2025; Hurst et al., 2024; Chen et al., 2024a), which excel at sequential data, to form and manipulate such spatial models remains a critical open question. Recent works have begun to probe this (Xu et al., 2025a; Zhang et al., 2024b; Chen et al., 2025c; Qi et al., 2025); for instance, MindCube (Liu et al., 2024) evaluates an agent’s ability to reason about a *static* 3D scene by reasoning about its complete layout from a few partial viewpoints. While also leveraging partial observations, CubeBench introduces the distinct challenge of long-horizon interaction with a *dynamic* cube. Our work therefore shifts the focus from reasoning about static perspectives to the more complex challenge of *updating* a mental model through direct, state-altering interaction with the environment, emphasizing mental simulation and long-horizon state tracking.

## G CONFIGURATION OF TEST SPLIT

Our evaluation set is constructed across eight distinct levels of task difficulty. These levels are determined by the state’s “depth”, which is the optimal number of moves required to solve a given cube configuration. For each of the eight difficulty levels, we sample five unique initial states. Each state is then tested across four different environment settings derived from the three-tiered framework: Full Symbolic (Tier 1), Full Visual (Tier 2), and the two Partial Visual modalities (Face View and Vertex View) from Tier 3. This methodology yields a total of 160 unique evaluation configurations (8 difficulties  $\times$  5 states  $\times$  4 settings) for each Large Language Model (LLM) agent. To provide an estimate of the resources required, a single, complete evaluation run of these 160 configurations on GPT-5 consumes a total of 59.3 million tokens (50.2 million input tokens and 9.1 million output tokens), resulting in a total cost of approximately \$153, based on the pricing of \$1.25 per million input tokens and \$10 per million output tokens.

## H INTRODUCTION TO THE POLICY GRADIENT BASELINE

As a baseline for comparison in the Full Symbolic setting, we implement a classic reinforcement learning agent based on the Policy Gradient (PG) method. Policy Gradient algorithms directly optimize the parameters of a policy by estimating the gradient of the expected return. The core idea is to adjust the policy’s parameters to increase the probability of taking actions that lead to higher

1026 cumulative rewards. Our implementation uses the REINFORCE algorithm, a foundational Monte  
 1027 Carlo policy gradient method.  
 1028

1029 The objective of the REINFORCE algorithm is to maximize the expected total discounted reward,  
 1030  $J(\theta)$ , by updating the policy parameters  $\theta$  in the direction of the gradient  $\nabla_{\theta}J(\theta)$ . The policy  
 1031 gradient theorem provides an estimate for this gradient:

$$1032 \nabla_{\theta}J(\theta) = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[ \sum_{t=0}^T \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) G_t \right]$$

$$1033$$

$$1034$$

1035 where  $\pi_{\theta}(a_t | s_t)$  is the policy (the probability of taking action  $a_t$  in state  $s_t$ ), and  $G_t = \sum_{k=t}^T \gamma^{k-t} r_k$   
 1036 is the total discounted return from time step  $t$  onward.  
 1037

1038 Our specific implementation utilizes a Multi-Layer Perceptron (MLP) to represent the policy net-  
 1039 work. The training is conducted using a curriculum learning strategy, where the agent is progres-  
 1040 sively trained on more difficult tasks by increasing the scramble length of the cube. The detailed  
 1041 configuration is as follows:

- 1042 • **Algorithm:** REINFORCE
- 1043 • **Policy Network:** A five-layer MLP with 256 neurons per layer and ReLU activation func-  
 1044 tions.
- 1045 • **Optimizer:** Adam with a learning rate of  $5 \times 10^{-4}$ .
- 1046 • **Discount Factor ( $\gamma$ ):** 0.99.
- 1047 • **Training Environment:** We use 64 parallel vectorized environments for efficient data col-  
 1048 lection.
- 1049 • **Update Rule:** The policy is updated after collecting a rollout of 512 steps from each par-  
 1050 allel environment, using a batch size of 512 episodes for the gradient update.
- 1051 • **Curriculum:** Training uses the curriculum learning for better convergence. We train two  
 1052 separate models, one for short-horizon tasks and one for long-horizon. For short-horizon  
 1053 tasks, the model is trained sequentially on scramble depths of 1, 2, 3, and 4, with the number  
 1054 of training timesteps scaled quadratically for each level (e.g., from 40,000 for depth 1 up  
 1055 to 320,000 for depth 4). For long-horizon tasks, a separate model is initialized from the  
 1056 converged short-horizon agent and then continues training on the more challenging depths  
 1057 from 5 to 20 with 320K steps each level.
- 1058 • **Max Number of Make-moves During Evaluation:** For short-horizon tasks, the maximum  
 1059 number of making moves is set to 16. For long-horizon ones, they are set to 400.

## 1062 I TOOLS FOR AGENTS

### 1064 I.1 FUNDAMENTAL INTERACTION TOOLS

1066 **make\_move:** This tool executes a single face rotation. It accepts one of 12 possible inputs corre-  
 1067 sponding to standard Singmaster notation: F (Front), B (Back), L (Left), R (Right), U (Up), and D  
 1068 (Down) for clockwise turns. A prime symbol ('') denotes a counter-clockwise rotation (e.g., F').  
 1069 Each call to this tool deterministically alters the cube's internal state. Since the agent can generate  
 1070 custom scripts in the `Code` block, this function can be called multiple times (e.g., in a loop) within  
 1071 a single step to execute a sequence of moves.

1072 **get\_observation:** This tool retrieves the current observation of the cube. The format of the  
 1073 returned data is contingent upon the experimental tier:

- 1075 • In **Tier 1 (Full Symbolic State)**, it returns a 54-character string that symbolically repre-  
 1076 sents the complete state of the cube.
- 1077 • In **Tier 2 (Full Visual State)** and **Tier 3 (Partial Visual State)**, it returns an image. De-  
 1078 pending on the specific task configuration, this can be a complete 2D unfolded map of the  
 1079 cube, an image of a single face (*face view*), or an image from a corner's perspective (*vertex  
 view*).

1080     **apply\_view\_transformation**: Exclusive to the **Tier 3 (Partial Visual State)** setting, this  
 1081     tool allows the agent to alter its viewpoint (e.g., up, down, left, right). This capability is essential for  
 1082     actively exploring the cube to reconstruct its full state from a series of limited views.  
 1083

1084     **I.2 AUXILIARY SOLVER TOOLS**  
 1085

1086     **StandardSolverTool** and **IdealSolverTool**: These tools provide the agent with access to  
 1087     a solver based on Kociemba’s two-phase algorithm (see Appendix K). They differ in a crucial aspect  
 1088     related to data formatting. The underlying solver requires a specific input format that is distinct from  
 1089     the state representation provided by the `get_observation` tool.

1090     


 1091         - The `StandardSolverTool` requires the agent to perform the necessary format conversion  
 1092             itself, thus testing its ability to transform data into a usable representation.
 1093         - The `IdealSolverTool` features a built-in converter. It allows the agent to input the  
 1094             state in the environment’s native format and receive a solution directly, thereby bypassing  
 1095             the format conversion challenge.

1096

1097     **J HEURISTIC ALGORITHM FOR SOLVING THE RUBIK’S CUBE:**  
 1098     **LAYER-BY-LAYER APPROACH**  
 1099

1100     This subsection formalizes the heuristic algorithm for solving the 3x3 Rubik’s Cube using the layer-  
 1101     by-layer (LBL) method. We drew on the method from *solvethecube* website<sup>2</sup> and adopted its illustrations.  
 1102     The cube-solving process is divided into seven major steps, each focusing on solving specific  
 1103     parts of the cube. Each step involves operations or algorithms that manipulate specific pieces, pro-  
 1104     gressively solving the puzzle. As introduced in our discussion on reward functions (Sec. 3.1.4),  
 1105     this LBL structure forms the basis of our Heuristic Metric ( $\phi_{heuristic}$ ). To calculate a score for any  
 1106     given cube state  $s$ , we evaluate it against the seven steps of the LBL method. The state’s score,  
 1107      $\phi_{heuristic}(s)$ , is defined as the highest step number (from 0 for a scrambled cube to 7 for a solved  
 1108     one) that the configuration has successfully completed.

1109     **Overall Strategy** The solving process begins with the initial scrambled state  $S_0$  and progresses  
 1110     through seven steps until the solved state  $S^*$  is achieved. Each transformation corresponds to a  
 1111     specific phase of the solution:

1112

$$S_0 \xrightarrow{f_1} S_1 \xrightarrow{f_2} S_2 \xrightarrow{f_3} \dots \xrightarrow{f_7} S^*$$

1113     where:

1114     


 1115         - $f_1$ : Forming the bottom cross.
 1116         - $f_2$ : Positioning the bottom corners.
 1117         - $f_3$ : Solving the second layer edges.
 1118         - $f_4$ : Creating the top cross.
 1119         - $f_5$ : Permuting the top edges.
 1120         - $f_6$ : Positioning the top corners.
 1121         - $f_7$ : Orienting the top corners.

1122     Each step is represented by an algorithm or set of moves that solves a specific portion of the cube  
 1123     without disrupting previously solved sections. See Fig. 7 for an intuitive schematic.

1124     **Step 1: Forming the Bottom Cross** The first task is to form a cross on the bottom layer by  
 1125     positioning the four edge pieces such that their colors match both the bottom face center and the  
 1126     adjacent side centers. This step can be formalized as:

1127     <sup>2</sup>URL:<https://solvethecube.com/>

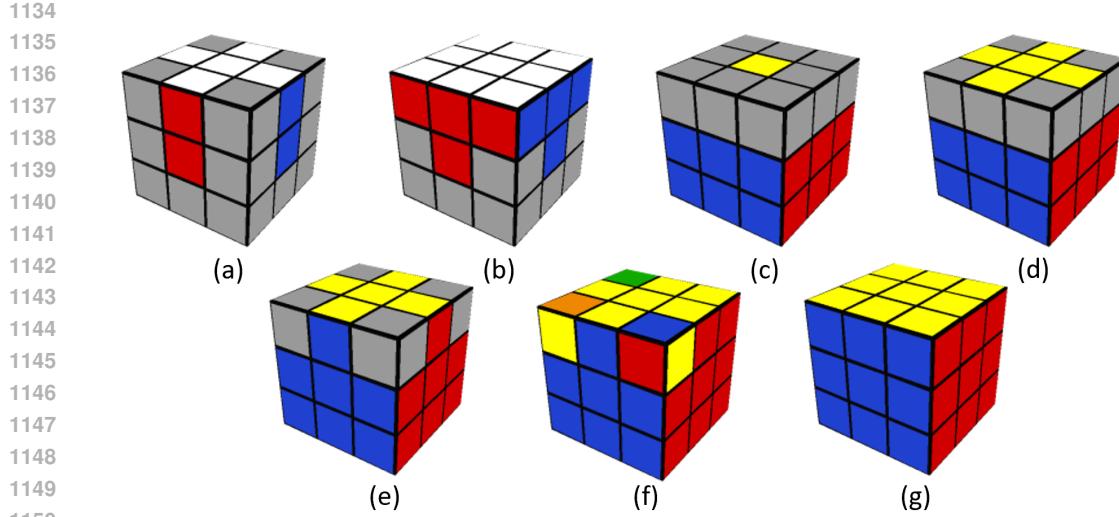


Figure 7: The diagram from solvethecube website illustrates a heuristic method for solving a Rubik's Cube: (a) Forming the bottom cross, (b) Positioning the bottom corners, (c) Solving the second layer edges, (d) Creating the top cross, (e) Permuting the top edges, (f) Positioning the top corners, and (g) Orienting the top corners.

$$f_{\text{cross}} : C_{\text{bottom}} \rightarrow C_{\text{bottom}}^*$$

where  $C_{\text{bottom}} = \{E_1, E_2, E_3, E_4\}$  represents the four edge pieces to be positioned, and  $C_{\text{bottom}}^*$  represents the state where the bottom cross is correctly formed. The process involves identifying the edge pieces and applying algorithms to move them into place.

**Step 2: Positioning the Bottom Corners** After the bottom cross is formed, the next objective is to position the bottom corner pieces. Let  $C_{\text{bottom-corner}} = \{C_1, C_2, C_3, C_4\}$  represent the four corner pieces to be positioned in the bottom layer. This step is formalized as:

$$f_{\text{corners}} : C_{\text{bottom-corner}} \rightarrow C_{\text{bottom-corner}}^*$$

where  $C_{\text{bottom-corner}}^*$  represents the correctly positioned bottom corners. The operation involves applying specific algorithms to move each corner piece into its correct location without disturbing the already solved bottom cross.

**Step 3: Solving the Second Layer Edges** The next step is to solve the edges of the second layer. Let  $C_{\text{second-layer-edge}} = \{E_5, E_6, E_7, E_8\}$  represent the four edge pieces that need to be positioned in the middle layer. The operation can be formalized as:

$$f_{\text{second-layer}} : C_{\text{second-layer-edge}} \rightarrow C_{\text{second-layer-edge}}^*$$

where  $C_{\text{second-layer-edge}}^*$  represents the state where the second layer edges are correctly positioned. The goal is to move the edge pieces from the top layer to the second layer, maintaining the solved bottom layer.

**Step 4: Creating the Top Cross** After solving the second layer, the next objective is to create a cross on the top layer. Let  $C_{\text{top-edge}} = \{E_9, E_{10}, E_{11}, E_{12}\}$  represent the four edge pieces that need to be positioned on the top layer. This step can be formalized as:

$$f_{\text{top-cross}} : C_{\text{top-edge}} \rightarrow C_{\text{top-edge}}^*$$

1188 where  $C_{\text{top-edge}}^*$  represents the state where the top layer edges are correctly aligned. The transformation  
 1189 for each edge piece is represented by:  
 1190

$$E_i \xrightarrow{\Delta_{E_i}} E_i^* \quad \text{for } i \in \{9, 10, 11, 12\}$$

1191 The goal is to position the edge pieces correctly on the top layer without disturbing the solved  
 1192 portions of the bottom and second layers.  
 1193

1194 **Step 5: Permuting the Top Edges** Once the top cross is formed, the next objective is to permute  
 1195 the top layer edge pieces into their correct positions. This operation can be formalized as:  
 1196

$$f_{\text{permute-edges}} : C_{\text{top-edge}}^* \rightarrow C_{\text{top-edge}}^{**}$$

1197 where  $C_{\text{top-edge}}^{**}$  represents the state where the top layer edges are correctly permuted. The goal is to  
 1198 apply specific algorithms that permute the top edges into their correct positions.  
 1199

1200 **Step 6: Positioning the Top Corners** After the top edges are permuted, the next step is to position  
 1201 the top layer corners. Let  $C_{\text{top-corner}} = \{C_{13}, C_{14}, C_{15}, C_{16}\}$  represent the four top corner pieces.  
 1202 This operation can be formalized as:  
 1203

$$f_{\text{position-corners}} : C_{\text{top-corner}} \rightarrow C_{\text{top-corner}}^*$$

1204 where  $C_{\text{top-corner}}^*$  represents the state where the top corners are positioned correctly. The goal is to  
 1205 apply specific algorithms to move the top layer corner pieces into their correct positions.  
 1206

1207 **Step 7: Orienting the Top Corners** The final step is to orient the top corners, ensuring that the  
 1208 top face becomes uniform in color. This operation can be formalized as:  
 1209

$$f_{\text{orient-corners}} : C_{\text{top-corner}}^* \rightarrow S^*$$

1210 where  $S^*$  represents the solved state of the Rubik's Cube. The transformation for each corner piece  
 1211 is represented by:  
 1212

$$C_i \xrightarrow{\Delta_{C_i}} C_i^* \quad \text{for } i \in \{13, 14, 15, 16\}$$

1213 The goal is to apply specific algorithms to orient the top corners without disturbing the already  
 1214 solved portions of the cube.  
 1215

1216 **Mathematical Summary of the Layer-by-Layer Approach** The Rubik's Cube solution can be  
 1217 mathematically summarized as a series of state transformations:  
 1218

$$S_0 \xrightarrow{f_{\text{cross}}} S_1 \xrightarrow{f_{\text{corners}}} S_2 \xrightarrow{f_{\text{second-layer}}} S_3 \xrightarrow{f_{\text{top-cross}}} S_4 \xrightarrow{f_{\text{permute-edges}}} S_5 \xrightarrow{f_{\text{position-corners}}} S_6 \xrightarrow{f_{\text{orient-corners}}} S^*$$

1219 Each transformation  $f_{\text{cross}}, f_{\text{corners}}, \dots, f_{\text{orient-corners}}$  corresponds to a specific set of moves that  
 1220 transform the cube toward the solved state. By applying these transformations in sequence, the cube is  
 1221 solved layer by layer.  
 1222

## 1223 K CUBE SOLVERS (TWO-PHASE AND OPTIMAL)

1224 This appendix introduces the two solvers used in our benchmark: the *Two-Phase Solver* and the *Optimal  
 1225 Solver*. The two-phase solver is used as the basic component for `StandardSolverTool`  
 1226 and `IdealSolverTool`, while the optimal solver is used for generating the testcases. In the following,  
 1227 we first outline the mathematical background for the two-phase method, and then describe  
 1228 the I/O formats and the roles each solver plays in our experiments.  
 1229

1242 K.1 INPUT AND OUTPUT FORMATS FOR CUBE STATE  
12431244 To interact with solvers, the cube state must be expressed in a precise *facelet string* representation:  
1245 a fixed ordering of all 54 stickers. Figure 8 illustrates the two indexing conventions used in our  
1246 system: (1) the **Initial Format** used by our environment, and (2) the **Solver Format** required by the  
1247 Kociemba two-phase solver.1248  
1249 **Initial Format.** The environment internally stores the cube as a 2D unfolded cross (see left of  
1250 Fig. 8). Each sticker is labeled by its face and index, e.g. U1–U9 for the Up face, R1–R9 for the  
1251 Right face, etc. Within each face, indices increase row by row from top-left to bottom-right:

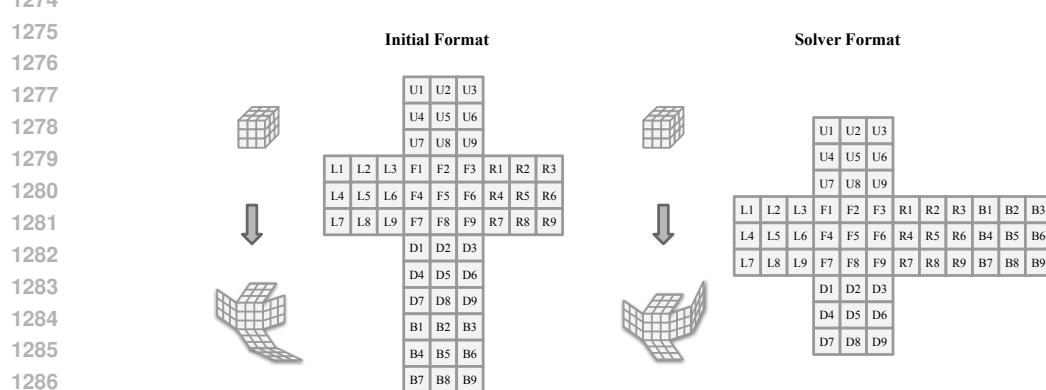
1252 
$$U1, U2, U3, U4, U5, U6, U7, U8, U9,$$
  
1253

1254 and similarly for  $R, F, D, L, B$ .1255 The environment’s *concatenation order* follows the visual “cross net” layout: first the Up face, then  
1256 Left–Front–Right in a row, followed by Down, and finally the Back face. Explicitly, the 54-character  
1257 string is constructed as

1258 
$$(F1, \dots, F9, B1, \dots, B9, L1, \dots, L9, R1, \dots, R9, U1, \dots, U9, D1, \dots, D9).$$
  
1259

1260 **Solver Format.** The two-phase solver requires the cube state as a 54-character string, concatenated  
1261 in the strict order

1262 
$$(U1, \dots, U9, R1, \dots, R9, F1, \dots, F9, D1, \dots, D9, L1, \dots, L9, B1, \dots, B9).$$
  
1263

1264 Each character encodes the color on the corresponding facelet. For example, the string UBL...  
1265 means: - position U1 has the U-color, - position U2 has the B-color, - position U3 has the L-color,  
1266 and so on. This flattened sequence is the standard *facelet string* convention used in Kociemba’s  
1267 solver.1268  
1269 **Conversion.** Figure 8 shows how the Initial Format is mapped into the Solver Format. This con-  
1270 version step is crucial: any misalignment (e.g., rotated faces or incorrect indexing) produces invalid  
1271 solver inputs and prevents the plan from being executed correctly. The Standard-Solver Agent must  
1272 handle the format conversion by itself, whereas the Ideal-Solver Agent has the conversion built in  
1273 `IdealSolverTool` and requires no further effort.1287  
1288 **Figure 8: Input/Output formats of cube state.** Left: Initial Format used in our environment. Right:  
1289 Right: Solver Format required by the two-phase solver. Both correspond to the same cube state but differ  
1290 in indexing layout and concatenation order.1291 K.2 TWO-PHASE SOLVER  
12931294 **Mathematical Background.** Let  $G_0$  denote the full Rubik’s Cube group generated by face turns  
1295 in the half-turn metric (HTM). Using standard Singmaster letters, a convenient presentation is

$$\begin{aligned}
1296 \quad G_0 &= \langle U, D, R, L, F, B \rangle \\
1297 &= \langle U, U', U^2, D, D', D^2, R, R', R^2, L, L', L^2, F, F', F^2, B, B', B^2 \rangle.
\end{aligned}$$

1299 That is, quarter turns, their inverses, and half turns generate the same group.

1300 The Two-Phase method (Kociemba; wik) reduces an arbitrary state to a structured subgroup and  
1301 then completes the solve inside that subgroup. Let  $H$  (often also written as  $G_1$ ) be the Phase-1  
1302 target subgroup,

$$\begin{aligned}
1304 \quad H &= \{ g \in G_0 \mid \text{all edges oriented, all corners oriented,} \\
1305 &\quad \text{and the four UD-slice edges lie in the UD slice} \}.
\end{aligned}$$

1306 Equivalently,  $H$  is exactly the set of states reachable using the restricted Phase-2 move set

$$\langle U, U', U^2, D, D', D^2, R^2, L^2, F^2, B^2 \rangle,$$

1309 since these moves preserve the above invariants. Finally,

$$\{e\} = G_2 \subset H \subset G_0,$$

1312 where  $e$  is the identity (the solved cube). The overall decomposition is

$$\text{Phase 1: } G_0 \rightarrow H \quad (\text{reduce to } H) \quad \text{Phase 2: } H \rightarrow \{e\} \quad (\text{solve within } H).$$

1315 In practice, both phases use depth-iterative search guided by large pruning (distance) tables; Phase 1  
1316 explores cosets of  $H$  in  $G_0$ , and Phase 2 searches within  $H$  down to  $e$ .

1318 **Algorithmic sketch.** Phase 1 finds a short maneuver (typically  $\leq 12$  HTM moves) that sends the  
1319 current state into  $H$ . Phase 2 continues from  $H$  using only  $U^{\pm 1}, U^2, D^{\pm 1}, D^2, R^2, L^2, F^2, B^2$  to  
1320 reach the identity. Due to strong heuristics and symmetry reductions, solutions are obtained quickly  
1321 and are usually  $\leq 20$  moves, though not formally guaranteed to be optimal (Kociemba).

1322 **Typical properties.** The method is fast, yields short, clean plans, and is well-suited as a callable  
1324 planning tool for agents.

### 1325 K.3 OPTIMAL SOLVER (IDA\* WITH COMPLETE PRUNING)

1327 **Algorithmic sketch.** The Optimal Solver performs Iterative Deepening A\* (IDA\*) search (Korf,  
1328 1985; 1997), guided by admissible heuristics derived from large pattern/pruning databases. IDA\*  
1329 combines the space-efficiency of depth-first search with the optimality guarantees of A\*. The search  
1330 depth limit is increased incrementally; once a solution is found at depth  $d$ , minimality is certified  
1331 since all shorter paths have been exhausted.

---

#### 1333 Algorithm 1 IDA\* with Pattern Database Heuristics

---

```

1: function IDA*( $s_0$ )
2:    $bound \leftarrow h(s_0)$ 
3:   while true do
4:      $t \leftarrow \text{SEARCH}(s_0, 0, bound)$ 
5:     if  $t = \infty$  then return failure
6:     else if  $t = \text{solution}$  then return solution path
7:     else
8:        $bound \leftarrow t$ 
9: function SEARCH( $s, g, bound$ )
10:   $f \leftarrow g + h(s)$ 
11:  if  $f > bound$  then return  $f$ 
12:  if  $s$  is goal then return solution
13:   $min \leftarrow \infty$ 
14:  for each successor  $s'$  of  $s$  do
15:     $t \leftarrow \text{SEARCH}(s', g + 1, bound)$ 
16:    if  $t = \text{solution}$  then return solution
17:    if  $t < min$  then
18:       $min \leftarrow t$ 
return  $min$ 

```

---

1350     **Heuristic.** The pruning tables (pattern databases) precompute exact solution lengths for subproblems such as subsets of edges or corners. During search, these values serve as admissible heuristics, 1351 drastically reducing the number of expanded states.  
1352

1353  
1354     **Typical properties.** Under HTM, the  $3 \times 3 \times 3$  cube can always be solved within 20 moves, 1355 with some positions requiring the exact 20 moves. (Rokicki et al., 2014; 2010). While run time 1356 can be seconds on easy instances, hard positions near the 20-move depth may require minutes to 1357 hours. However, the returned solution is provably optimal, making the solver suitable for dataset 1358 construction and difficulty certification.  
1359

1360     Table 7: Two-Phase vs. Optimal Solvers.  
1361

Solver	Optimality	Role
Two-Phase	Near-optimal	Agent tool
Optimal	Provably optimal	Testcase construction

1366     L CASE STUDIES  
1367

- 1370     • Sec. L.1: An agent solves the cube by applying a heuristic beam search.
- 1371     • Sec. L.2: An agent tries first with IDDFS and then with the Beginner’s Method.
- 1372     • Sec. L.3: A less capable agent attempts a random walk (brute-force) and fails to make progress.
- 1373     • Sec. L.4: An agent implements a meet-in-the-middle search, but its pruning strategies are insufficient to overcome the exponential state space.
- 1374     • Sec. L.5: An agent bypasses visual reasoning, relying exclusively on a heuristic reward to guide 1375 its symbolic search.
- 1376     • Sec. L.6: A capable agent is misled by a simplistic reward and only succeeds after abandoning it 1377 in favor of its own subgoal-based reasoning.
- 1378     • Sec. L.7: A simple external reward successfully guides an agent’s search, leading to incremental 1379 progress.
- 1380     • Sec. L.8: An agent learns to use an external planner by systematically experimenting with input 1381 formats.
- 1382     • Sec. L.9: An agent succeeds on a structured visual task by using a symbolic shortcut but fails on 1383 a distorted view that requires genuine spatial reasoning.

1384     L.1 BASIC AGENT – BEAM SEARCH  
1385

Model	Agent Type	Input	Reward	Depth
GPT-5	Basic	Full Symbolic	No reward	8

1386     In this case, GPT-5 applies a heuristic beam search strategy to solve the Rubik’s Cube. Initially, the 1387 agent deduces the exact permutations of the six basic moves (F, B, U, D, R, L) using randomized 1388 scrambles and signature matching. With these move permutations, the agent can simulate cube state 1389 transitions internally, allowing it to conduct search without repeated environment queries.  
1390

1391     The search itself is implemented as a beam-style search with heuristic pruning:  
1392

1393     **Heuristic function:** The agent uses a simple facelet mismatch count to evaluate each state relative 1394 to the solved configuration.  
1395

1396     **Beam width:** At each depth, candidate states are sorted by heuristic value, and only the top 1,500 1397 candidates are retained (`frontier = candidates[:beam_width]`), limiting the branching 1398 factor.  
1399

1404     **Pruning rules:** Immediate inverse moves are skipped (if `last` and `inverse[last] ==`  
 1405     `mv`: `continue`), and already visited states are not revisited (if `new_state` in `visited`:  
 1406     `continue`).  
 1407

1408     During search, the frontier grows progressively with depth (e.g., 12 → 114 → 1,068 candidates  
 1409     in the first three layers), demonstrating layered expansion with controlled breadth. Eventually, the  
 1410     agent identifies a four-move solution, applies it, and verifies that the cube reaches the solved state.

1411     **Step 3**

1412     **Thought.**

1414  
 1415     Thought: The previous attempt failed because the assumed face rotation orientations for U, D, R, L did  
 1416     not match the environment. I will now robustly deduce the exact permutation for each basic move (F, B,  
 1417     U, D, R, L) by signature matching: I'll run multiple randomized scrambles, snapshot the state before  
 1418     and after applying the move, and then reconstruct the index permutation by matching the K-length color  
 1419     signature for each index. This is independent of face orientations. With all six move permutations  
 1420     learned, I'll perform a heuristic beam search offline to find a solution, then apply the move sequence  
 1421     with `make_move` and verify. If the cube becomes solved, I'll call `final_answer("")`.

1422     **Code.**

```
1424
1425 import random
1426 import time
1427
1428 # Utilities
1429 def inverse_move(mv):
1430     return mv[:-1] if mv.endswith(" ") else mv + " "
1431
1432 def apply_seq(seq):
1433     for m in seq:
1434         make_move(m)
1435
1436 def invert_seq(seq):
1437     return [inverse_move(m) for m in reversed(seq)]
1438
1439 def deduce_move_by_signature(move, trials=7, scramble_len=12):
1440     # Ensure we start from a known base and return to it after each trial and at the end.
1441     base0 = get_observation()
1442     pre_list = []
1443     post_list = []
1444
1445     # Allowed moves for scrambling
1446     moves = ['F', 'B', 'U', 'D', 'R', 'L']
1447     all_moves = moves + [m + " " for m in moves]
1448
1449     for t in range(trials):
1450         # Make a scramble that avoids immediate inverses
1451         scramble = []
1452         last = None
1453         for k in range(scramble_len):
1454             while True:
1455                 m = random.choice(all_moves)
1456                 if last and inverse_move(last) == m:
1457                     continue
1458                 scramble.append(m)
1459                 last = m
1460                 break
1461
1462             # Apply scramble
1463             apply_seq(scramble)
1464             s_pre = get_observation()
1465             # Apply the move to learn
1466             make_move(move)
1467             s_post = get_observation()
1468             # Undo the move
1469             make_move(inverse_move(move))
1470             # Undo scramble
1471             apply_seq(invert_seq(scramble))
1472             # Verify we returned to base0
1473             if get_observation() != base0:
1474                 print("Warning: state mismatch after undoing scramble during deduction of", move)
1475                 pre_list.append(s_pre)
1476                 post_list.append(s_post)
```

```

1458
1459     # Build signature vectors for each index
1460     pre_sigs = [".".join(pre[i] for pre in pre_list) for i in range(54)]
1461     post_sigs = [".".join(post[i] for post in post_list) for i in range(54)]
1462
1463     # Map post index j to pre index i such that signatures match
1464     perm = [None]*54
1465     # Build dict from signature to indices for pre
1466     sig_to_indices = {}
1467     for i, sig in enumerate(pre_sigs):
1468         sig_to_indices.setdefault(sig, []).append(i)
1469
1470     ambiguous = []
1471     missing = []
1472     for j, sig in enumerate(post_sigs):
1473         candidates = sig_to_indices.get(sig, [])
1474         if len(candidates) == 1:
1475             perm[j] = candidates[0]
1476         elif len(candidates) == 0:
1477             missing.append(j)
1478         else:
1479             ambiguous.append((j, candidates))
1480
1481     if missing or ambiguous:
1482         # Try additional trials to disambiguate
1483         extra = 0
1484         max_extra = 6
1485         while (missing or ambiguous) and extra < max_extra:
1486             extra += 1
1487             # Do one extra trial
1488             scramble = []
1489             last = None
1490             for k in range(scramble_len):
1491                 while True:
1492                     m = random.choice(all_moves)
1493                     if last and inverse_move(last) == m:
1494                         continue
1495                     scramble.append(m)
1496                     last = m
1497                     break
1498             apply_seq(scramble)
1499             s_pre = get_observation()
1500             make_move(move)
1501             s_post = get_observation()
1502             make_move(inverse_move(move))
1503             apply_seq(invert_seq(scramble))
1504             pre_list.append(s_pre)
1505             post_list.append(s_post)
1506
1507             # Recompute signatures
1508             pre_sigs = [".".join(pre[i] for pre in pre_list) for i in range(54)]
1509             post_sigs = [".".join(post[i] for post in post_list) for i in range(54)]
1510             sig_to_indices = {}
1511             for i, sig in enumerate(pre_sigs):
1512                 sig_to_indices.setdefault(sig, []).append(i)
1513             perm = [None]*54
1514             missing = []
1515             ambiguous = []
1516             for j, sig in enumerate(post_sigs):
1517                 candidates = sig_to_indices.get(sig, [])
1518                 if len(candidates) == 1:
1519                     perm[j] = candidates[0]
1520                 elif len(candidates) == 0:
1521                     missing.append(j)
1522                 else:
1523                     ambiguous.append((j, candidates))
1524
1525             if missing or ambiguous:
1526                 print(f"Could not uniquely deduce mapping for {move}. Missing={len(missing)}, ambiguous={len(ambiguous)}")
1527             return None
1528
1529     # Validate by applying to a fresh snapshot
1530     base = get_observation()
1531     make_move(move)
1532     s1 = get_observation()
1533     make_move(inverse_move(move))
1534
1535     # Build simulated application
1536     sim = "".join(base[perm[i]] for i in range(54))
1537     if sim != s1:
1538         print(f"Validation failed for move {move}.")
1539         return None
1540
1541     return perm
1542
1543
1544     # Learn all base move permutations

```

```

1512
1513     moves = ['F', 'B', 'U', 'D', 'R', 'L']
1514     perm_map = {}
1515     for mv in moves:
1516         print("Learning permutation for", mv)
1517         p = deduce_move_by_signature(mv, trials=7, scramble_len=10)
1518         if p is None:
1519             print("Failed to learn move:", mv)
1520             # Try a second attempt with more trials if failed
1521             p = deduce_move_by_signature(mv, trials=10, scramble_len=12)
1522         if p is None:
1523             print("Giving up on move:", mv)
1524             break
1525     perm_map[mv] = p
1526
1527     if len(perm_map) != 6:
1528         print("Failed to learn all permutations. Learned:", list(perm_map.keys()))
1529     else:
1530         # Build prime move perms
1531         def invert_perm(perm):
1532             inv = [0]*54
1533             for i, src in enumerate(perm):
1534                 inv[src] = i
1535             return inv
1536         for mv in moves:
1537             perm_map[mv+"'"] = invert_perm(perm_map[mv])
1538
1539         def sim_apply_state(state, move):
1540             perm = perm_map[move]
1541             return "".join(state[perm[i]] for i in range(54))
1542
1543         # Solved target
1544         solved = "R"*9 + "G"*9 + "B"*9 + "Y"*9 + "O"*9 + "W"*9
1545
1546         def heuristic(state):
1547             # Simple facelet mismatch
1548             return sum(1 for i, ch in enumerate(state) if ch != solved[i])
1549
1550         # Beam search
1551         all_moves = moves + [m+"'" for m in moves]
1552         inverse = {m: m+"'" for m in moves}
1553         inverse.update({m+"'": m for m in moves})
1554
1555         start = get_observation()
1556         if start == solved:
1557             final_answer("")
1558         else:
1559             beam_width = 1500
1560             max_depth = 20
1561
1562             frontier = [(start, [])]
1563             visited = set([start])
1564             best_h = heuristic(start)
1565             print("Initial heuristic:", best_h)
1566
1567             solved_path = None
1568
1569             for depth in range(1, max_depth+1):
1570                 candidates = []
1571                 for state, path in frontier:
1572                     last = path[-1] if path else None
1573                     for mv in all_moves:
1574                         if last and inverse[last] == mv:
1575                             continue
1576                         new_state = sim_apply_state(state, mv)
1577                         if new_state in visited:
1578                             continue
1579                         visited.add(new_state)
1580                         new_path = path + [mv]
1581                         if new_state == solved:
1582                             solved_path = new_path
1583                             break
1584                         h = heuristic(new_state)
1585                         candidates.append((h, new_state, new_path))
1586
1587             if solved_path is not None:
1588                 break
1589
1590             if solved_path is not None:
1591                 print("Found solution of length", len(solved_path))
1592                 break
1593             candidates.sort(key=lambda x: (x[0], len(x[2])))
1594             frontier = [(st, p) for (h, st, p) in candidates[:beam_width]]
1595             print(f"Depth {depth}: frontier={len(frontier)}, best_h_so_far={min([h for (h, _, _) in candidates], default=best_h)}")
1596
1597             if not frontier:

```

```

1566
1567         break
1568
1569     if solved_path is None:
1570         print("Beam search did not find a solution.")
1571     else:
1572         print("Applying solution:", " ".join(solved_path))
1573         apply_seq(solved_path)
1574         final_state = get_observation()
1575         print("Final state:", final_state)
1576         if final_state == solved:
1577             final_answer(" ")
1578         else:
1579             print("Solution application did not reach solved state. Mismatches:",
1580             heuristic(final_state))
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619

```

### Observation.

```

Learning permutation for F
Learning permutation for B
Learning permutation for U
Learning permutation for D
Learning permutation for R
Learning permutation for L
Initial heuristic: 36
Depth 1: frontier=12, best_h_so_far=31
Depth 2: frontier=114, best_h_so_far=22
Depth 3: frontier=1068, best_h_so_far=12
Found solution of length 4
Applying solution: B R' F U
Final state: RRRRRRRRRGGGGGGGGBBBBBBBBYYYYYYYYYOOOOOOOOOWWWWWWWWW

```

*Solved.*

## L.2 BASIC AGENT – IDDFS AND BEGINNER’S METHOD

Model	Agent Type	Input	Reward	Depth
gpt-oss-120b	Basic	Full Symbolic	No reward	8

In this case, the agent first performs a depth-limited iterative deepening search `dfs(depth, max_depth, last_face)` systematically explores move sequences, applying each candidate move with `apply(m)` and backtracking via `undo(m)`. Pruning strategies are directly reflected in the code: `if last_face is not None and face_of(m) == last_face: continue` prevents consecutive moves on the same face, and `is_inverse(a, b)` is used to avoid immediate inverse moves. Early goal checks (`if is_solved_state(get_observation())`) terminate branches once a solution is found. Iterative deepening is implemented by looping over depth limits (`for d in range(limit_depth+1)`).

Unfortunately, several searches failed. After a few failures, it attempted a different strategy.

In its thoughts (Steps 8-10), the agent correctly outlines the sub-goals of the beginner’s method, from building the white cross to solving the final layer with OLL (Orientation of the Last Layer) and PLL (Permutation of the Last Layer) algorithms.

However, the agent’s execution is undermined by a subtle but critical flaw in its internal world model—an incorrect mapping of the cube’s indices to its symbolic representation. The agent itself explicitly identifies this error in a moment of self-correction in Step 13, noting that “earlier mappings had mistakes.” Despite this high-level awareness, the cognitive load of implementing the entire long-horizon plan within a single, complex code block introduces new, unperceived errors. The final outcome is a failure, where the cube’s state is no closer to being solved.

*Step 4*

**Thought.**

1620  
1621  
1622  
1623  
1624  
1625  
1626  
1627  
1628  
1629

Thought: The previous searches were too deep or inefficient. I will attempt a carefully pruned iterative deepening search up to depth 5 using a reduced move set that includes quarter and half turns, with early goal checks at every node and strict pruning (no repeating the same face, and no immediate inverses). This should keep the number of operations manageable. If a solution exists within 5 moves, it will be found and applied directly on the real cube; then I will call final\_answer("") .

## Code.

```

# Depth-limited IDDFS with pruning and early goal checks.
# Uses moves including half turns (implemented as two quarter turns).

from time import time

def is_solved_state(s):
    return (all(c == 'R' for c in s[0:9]) and
            all(c == 'G' for c in s[9:18]) and
            all(c == 'B' for c in s[18:27]) and
            all(c == 'Y' for c in s[27:36]) and
            all(c == 'O' for c in s[36:45]) and
            all(c == 'W' for c in s[45:54]))

start = get_observation()
if is_solved_state(start):
    final_answer("")
else:
    # Define move variants with half-turns
    basic_faces = ['F', 'B', 'L', 'R', 'U', 'D']
    moves = []
    for f in basic_faces:
        moves.append(f)
        moves.append(f+"1")
        moves.append(f+"2")
    # Helper to apply and undo moves in the environment
    def apply(m):
        if m.endswith("2"):
            face = m[0]
            make_move(face)
            make_move(face)
        else:
            make_move(m)
    def undo(m):
        if m.endswith("2"):
            face = m[0]
            make_move(face)
            make_move(face)
        else:
            # inverse of X is X', inverse of X' is X
            if m.endswith("1"):
                make_move(m[0])
            else:
                make_move(m+"1")
    def face_of(m):
        return m[0]
    def is_inverse(a, b):
        # Treat 2-turn as self-inverse
        if a[0] != b[0]:
            return False
        if a.endswith("2") and b.endswith("2"):
            return True
        if a.endswith("1") and (b == a[0]):
            return True
        if (a == a[0]) and b.endswith("1") and (b[0] == a[0]):
            return True
        return False

    max_nodes = 300000 # safeguard against excessive search
    node_count = 0
    found = False
    solution_path = []

    def dfs(depth, max_depth, last_face):
        nonlocal node_count, found, solution_path
        if found:
            return True
        node_count += 1
        if node_count > max_nodes:
            return False

```

```

1674
1675     # Early goal check at every node
1676     if is_solved_state(get_observation()):
1677         found = True
1678         return True
1679     if depth == max_depth:
1680         return False
1681     for m in moves:
1682         # pruning: avoid repeating the same face consecutively
1683         if last_face is not None and face_of(m) == last_face:
1684             continue
1685         # pruning: avoid immediate inverse with previous move in path if exists
1686         prev = solution_path[-1] if solution_path else None # note: we append on success; here
1687         # Apply move
1688         apply(m)
1689         if dfs(depth+1, max_depth, face_of(m)):
1690             solution_path.append(m)
1691             return True
1692         # Backtrack
1693         undo(m)
1694     return False
1695
1696     # Iterative deepening up to depth 5
1697     start_time = time()
1698     limit_depth = 5
1699     for d in range(limit_depth+1):
1700         if dfs(0, d, None):
1701             break
1702     # If found, the moves leading to solution were applied in-place during recursion.
1703     if found and is_solved_state(get_observation()):
1704         final_answer("")
1705     else:
1706         print(f"Search exhausted (nodes={node_count}). Could not solve within depth {limit_depth}.")

```

## Observation.

```

1701 <Killed by the Python interpreter>
1702 Maybe there is an infinite loop somewhere in the code, or you're just asking too many calculations.

```

Step 5, 6 still failed. Changed strategy at Step 7.

## Thought.

```

1709 Thought: Our previous brute-force search approaches hit interpreter limits. I will now implement a
1710 deterministic, step-by-step white-cross builder (beginner's method) that:
1711 - Analyzes the current state string (observation-type is state_string).
1712 - Locates each white edge (W with its partner color).
1713 - Applies targeted moves to:
1714     - Lift the white edge to the U layer with white facing up using single face turns.
1715     - Rotate U (and D when needed) to align the edge above its correct face.
1716     - Insert with a double face turn (X2) to place it correctly in the D (white) cross while keeping
1717     previously placed edges intact.
1718
1719 I will implement robust mapping for edges and use only simple, safe moves (F/F'/F2, R/R'/R2, B/B'/B2,
1720 L/L'/L2, U/U', D/D'). I will repeat for the four target white edges: WR (to FD), WY (to RD), WG (to
1721 BD), WB (to LD). After building the white cross, I will check if the cube is already solved; if it is,
1722 I will call final_answer(""). If not, we can proceed further in subsequent steps (but for now, try to
1723 get the cross in place). Let's execute this first stage.

```

## Code.

```

1724 # Build the white cross step-by-step using targeted moves.
1725 # Faces and centers per the problem:
1726 #   F (front) = Red (R)
1727 #   B (back)   = Green (G)
1728 #   L (left)   = Blue (B)
1729 #   R (right)  = Yellow (Y)

```

```

1728
1729     #   U (up)      = Orange (O)
1730     #   D (down)    = White (W)
1731
1732     # Helper: Apply a move, double move, inverse move
1733     def do(m):
1734         make_move(m)
1735
1736     def do2(face):
1737         make_move(face); make_move(face)
1738
1739     def inv(m):
1740         return m[:-1] if m.endswith("") else (m + "")
1741
1742     def is_solved_state(s):
1743         return (all(c == 'R' for c in s[0:9]) and
1744                 all(c == 'G' for c in s[9:18]) and
1745                 all(c == 'B' for c in s[18:27]) and
1746                 all(c == 'Y' for c in s[27:36]) and
1747                 all(c == 'O' for c in s[36:45]) and
1748                 all(c == 'W' for c in s[45:54]))
1749
1750     # Centers indices
1751     def centers_of(state):
1752         return {
1753             'F': state[4],    # should be 'R'
1754             'B': state[13],  # should be 'G'
1755             'L': state[22],  # should be 'B'
1756             'R': state[31],  # should be 'Y'
1757             'U': state[40],  # should be 'O'
1758             'D': state[49],  # should be 'W'
1759         }
1760
1761     # Edge map: name -> ((face1, idx1), (face2, idx2))
1762     edge_defs = {
1763         'UF': (('U', 42), ('F', 1)),
1764         'UR': (('U', 41), ('R', 28)),
1765         'UB': (('U', 37), ('B', 10)),
1766         'UL': (('U', 39), ('L', 19)),
1767         'FR': (('F', 5), ('R', 30)),
1768         'FL': (('F', 3), ('L', 23)),
1769         'FD': (('F', 7), ('D', 46)),
1770         'RB': (('R', 32), ('B', 12)),
1771         'RD': (('R', 34), ('D', 50)),
1772         'LB': (('L', 21), ('B', 14)),
1773         'LD': (('L', 25), ('D', 48)),
1774         'BD': (('B', 16), ('D', 52)),
1775     }
1776
1777     # For each face's "UX" and "XD" reference
1778     UX_of = {'F': 'UF', 'R': 'UR', 'B': 'UB', 'L': 'UL'}
1779     XD_of = {'F': 'FD', 'R': 'RD', 'B': 'BD', 'L': 'LD'}
1780
1781     # Rotations of U that cycle UF->UR->UB->UL->UF and likewise the edge names
1782     U_cycle = ['UF', 'UR', 'UB', 'UL']
1783     # Rotations of D that cycle FD->LD->BD->RD->FD under D' (counterclockwise) or FD->RD->BD->LD->FD under D
1784     D_cycle_clockwise = ['FD', 'RD', 'BD', 'LD']
1785
1786     # Helper to read current state char at a given absolute index
1787     def read(state, idx):
1788         return state[idx]
1789
1790     def get_edge_colors(state, edge_name):
1791         (f1, i1), (f2, i2) = edge_defs[edge_name]
1792         return (state[i1], state[i2])
1793
1794     def locate_white_edge_partner(state, partner_color):
1795         # Returns (edge_name, where_white) where where_white in {0,1} indexes which of the two stickers is
1796         # white
1797         for name, ((f1, i1), (f2, i2)) in edge_defs.items():
1798             c1, c2 = state[i1], state[i2]
1799             if (c1 == 'W' and c2 == partner_color):
1800                 return name, 0
1801             if (c2 == 'W' and c1 == partner_color):
1802                 return name, 1
1803         return None, None
1804
1805     # Rotate U k times (k in {0..3})
1806     def U_times(k):
1807         for _ in range(k % 4):
1808             do('U')
1809
1810     # Rotate D k times clockwise

```

```

1782
1783     def D_times(k):
1784         for _ in range(k % 4):
1785             do('D')
1786
1787         # Align a U-layer edge to be at a specific UX position by rotating U
1788         def align_U_edge_to(state, edge_name_current, target_UX):
1789             # edge_name_current must be one of U-layer edges: 'UF', 'UR', 'UB', 'UL'
1790             # compute offset
1791             cur_idx = U_cycle.index(edge_name_current)
1792             tgt_idx = U_cycle_clockwise.index(target_UX)
1793             k = (tgt_idx - cur_idx) % 4
1794             U_times(k)
1795
1796         # Align a D-layer edge to be at a specific XD position by rotating D
1797         def align_D_edge_to(state, edge_name_current, target_XD):
1798             cur_idx = D_cycle_clockwise.index(edge_name_current)
1799             tgt_idx = D_cycle_clockwise.index(target_XD)
1800             k = (tgt_idx - cur_idx) % 4
1801             D_times(k)
1802
1803         # Given an edge in the middle layer, lift it to U with white sticker on U using a single move.
1804         def lift_middle_white_to_U(state, edge_name, white_on_face):
1805             # white_on_face is 'F', 'R', 'B', 'L' representing which face holds the white sticker for this edge.
1806             # For each middle edge we choose the face turn that moves it to the U layer with white on U.
1807             if edge_name == 'FR':
1808                 if white_on_face == 'F':
1809                     do("F")
1810                 elif white_on_face == 'R':
1811                     do("R")
1812             elif edge_name == 'FL':
1813                 if white_on_face == 'F':
1814                     do("F")
1815                 elif white_on_face == 'L':
1816                     do("L")
1817             elif edge_name == 'RB':
1818                 if white_on_face == 'R':
1819                     do("R")
1820                 elif white_on_face == 'B':
1821                     do("B")
1822             elif edge_name == 'LB':
1823                 if white_on_face == 'L':
1824                     do("L")
1825                 elif white_on_face == 'B':
1826                     do("B")
1827             else:
1828                 # Not a middle edge
1829                 pass
1830
1831         # If in U-layer but white on side, drop it to D with white on D using a single face turn.
1832         def drop_U_side_white_to_D(state, edge_name, white_on_face):
1833             # edge_name one of UF, UR, UB, UL; white_on_face in ('F', 'R', 'B', 'L')
1834             if edge_name == 'UF' and white_on_face == 'F':
1835                 do("F")
1836             elif edge_name == 'UR' and white_on_face == 'R':
1837                 do("R")
1838             elif edge_name == 'UB' and white_on_face == 'B':
1839                 do("B")
1840             elif edge_name == 'UL' and white_on_face == 'L':
1841                 do("L")
1842
1843         # Eject from D when white on side: bring under that face then do that face' to bring to middle/U
1844         def eject_D_side_white(state, edge_name, white_on_face):
1845             # edge_name in ('FD', 'RD', 'BD', 'LD'), white_on_face in ('F', 'R', 'B', 'L')
1846             # First rotate D to bring it under the face that has the white sticker
1847             target = white_on_face + 'D'
1848             align_D_edge_to(state, edge_name, target)
1849             # Now do the inverse face to move the edge up into the middle layer where we can lift it in next
1850             steps
1851             do(white_on_face + "")
1852
1853         # After white edge is on U face with white up at some UX, align above its target face and insert with
1854         # X2
1855         def insert_from_U_white_up(state, target_face):
1856             # target_face in {'F', 'R', 'B', 'L'}
1857             # Bring the white edge to UX position by U-turns, then do target_face
1858             # To know current UX position, find WR edge in U layer:
1859             # But we call this only after aligning with previous helpers; safe to rotate U until the partner
1860             # color aligns with target center
1861             # We'll explicitly rotate U until the edge is at UX.
1862             # Determine which UX currently holds the (W, partner) edge
1863             desired_edge_name = UX_of[target_face]
1864             # Find the current U-edge location of the piece with colors {W, centers[target_face]}
1865             state2 = get_observation()

```

```

1836
1837     centers = centers_of(state2)
1838     partner = centers[target_face]
1839     # Find among U-edges where this WR is
1840     u_positions = ['UF', 'UR', 'UB', 'UL']
1841     current_u = None
1842     where_white = None
1843     for en in u_positions:
1844         (f1, i1), (f2, i2) = edge_defs[en]
1845         c1, c2 = state2[i1], state2[i2]
1846         if (c1 == 'W' and c2 == partner):
1847             current_u, where_white = en, f1 # white on f1
1848             break
1849         if (c2 == 'W' and c1 == partner):
1850             current_u, where_white = en, f2
1851             break
1852     if current_u is None:
1853         return
1854     # Rotate U to bring it to UX_of[target_face]
1855     align_U_edge_to(state2, current_u, desired_edge_name)
1856     # Perform the double turn to insert
1857     do2(target_face)

1858     # Main white cross builder
1859     def build_white_cross():
1860         # Process targets in order: F (R), R (Y), B (G), L (B)
1861         order = ['F', 'R', 'B', 'L']
1862         max_outer_iters = 80
1863         for face in order:
1864             iters = 0
1865             while iters < 40:
1866                 iters += 1
1867                 st = get_observation()
1868                 ctr = centers_of(st)
1869                 partner = ctr[face] # desired partner color for this edge
1870                 # Check if it's already solved at XD with correct orientation
1871                 xd = XD_of[face]
1872                 (f1, i1), (f2, i2) = edge_defs[xd]
1873                 c1, c2 = st[i1], st[i2]
1874                 # Determine orientation at XD: white on D position?
1875                 if (c1 == partner and c2 == 'W') or (c2 == partner and c1 == 'W'):
1876                     # Additionally ensure white is on D (i.e., the index corresponding to D in this XD)
1877                     # For XD edge, the pair is (face, idx_face), (D, idx_down). The D index may be either
1878                     # i1 or i2.
1879                     # Identify which is D:
1880                     down_idx = i1 if f1 == 'D' else (i2 if f2 == 'D' else None)
1881                     if down_idx is not None and st[down_idx] == 'W':
1882                         # Solved for this edge
1883                         break
1884
1885                     # Otherwise, locate this white-partner edge
1886                     edge_name, pos = locate_white_edge_partner(st, partner)
1887                     if edge_name is None:
1888                         # Should not happen
1889                         break
1890
1891                     # Determine which faces the two stickers are on and their colors
1892                     (ef1, ei1), (ef2, ei2) = edge_defs[edge_name]
1893                     col1, col2 = st[ei1], st[ei2]
1894
1895                     # Case handling
1896                     if edge_name in ['UF', 'UR', 'UB', 'UL']:
1897                         # U layer
1898                         # Identify where white is: on U or on side
1899                         white_on = None
1900                         if st[ei1] == 'W':
1901                             white_on = ef1
1902                         elif st[ei2] == 'W':
1903                             white_on = ef2
1904                         if white_on == 'U':
1905                             # Rotate U to UX and insert with face2
1906                             insert_from_U_white_up(st, face)
1907                         else:
1908                             # White on side face in U layer: drop to D with single move on that face
1909                             drop_U_side_white_to_D(st, edge_name, white_on)
1910                     elif edge_name in ['FR', 'FL', 'RB', 'LB']:
1911                         # Middle layer: lift to U with white on U
1912                         white_on = None
1913                         if st[ei1] == 'W':
1914                             white_on = ef1
1915                         elif st[ei2] == 'W':
1916                             white_on = ef2
1917                         lift_middle_white_to_U(st, edge_name, white_on)
1918                     elif edge_name in ['FD', 'RD', 'BD', 'LD']:
1919                         # Bottom layer: drop to D with single move on that face
1920                         drop_B_bottom_white_to_D(st, edge_name, white_on)
1921
1922     return state2

```

```

1890
1891     # D layer
1892     # Determine if white is on D or on side
1893     white_on = None
1894     if st[ei1] == 'W':
1895         white_on = efl
1896     elif st[ei2] == 'W':
1897         white_on = efr
1898     if white_on == 'D':
1899         # Rotate D to bring under target face, insert with face2
1900         current_xd = edge_name
1901         align_D_edge_to(st, current_xd, XD_of[face])
1902         do2(face)
1903     else:
1904         # White on side at D layer: eject up to middle/U
1905         eject_D_side_white(st, edge_name, white_on)
1906     else:
1907         # Should not happen
1908         pass
1909     # Next target
1910
1911 # End for faces
1912
1913 # First, if solved already, finish
1914 st0 = get_observation()
1915 if is_solved_state(st0):
1916     final_answer("")
1917 else:
1918     # Build white cross
1919     build_white_cross()
1920     # Check if solved by chance, else continue later
1921     st1 = get_observation()
1922     if is_solved_state(st1):
1923         final_answer("")
1924     else:
1925         # Print progress: show current state to proceed next steps in subsequent iteration
1926         print(get_observation())
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943

```

## Observation.

```
OGGYRRRRBBORWRGBWGGYWBYBGYOBWOYOYGOYGOYGROOBOYGWBWRRBWWRWY
```

## Step 8

### Thought.

I will continue solving using the beginner's method. I already built some progress toward the white cross. Next, I will insert the four white corners into the D (white) layer without disturbing the cross, using standard triggers:

- For DFR (from UFR): R U R' U'
- For DRB (from URB): B U B' U'
- For DBL (from UBL): L' U' L U
- For DLF (from ULF): F' U' F U

I will:

- Define corner indices.
- Detect each target white corner by its color set {W, center(face1), center(face2)}.
- If found in D but unsolved: rotate D to bring it to the correct slot and eject with the corresponding trigger until it pops to U.
- If found in U: rotate U to place above its target slot, then apply the corresponding trigger repeatedly until solved.
- Repeat for all four corners, then print the current state. If solved, I will call final\_answer("") .

## Step 9

### Thought.

I will proceed with the next stage of the beginner's method: solve the middle layer edges (second layer) without disturbing the completed white layer. I will:

- Identify middle layer target edges (FR, FL, RB, LB).

```

1944
1945     - Insert correct U-layer edges that do not contain U or D colors (i.e., not containing O or W).
1946     - If no such U edges exist, eject an incorrectly placed middle edge to the U layer, then insert it
1947     correctly.
1948     I will then print the updated cube state (or call final_answer("") if by chance the cube becomes
1949     solved).
1950
1951
1952 Step 10
1953 Thought.
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997

```

The cube now has the first two layers solved. I will complete the last layer using beginner's method:  
 - OLL: orient U (Orange) edges to a cross using the FRUR'U'F' sequence with correct U rotations; then  
 orient corners using Sune/anti-Sune until all U stickers are Orange.  
 - PLL: first permute last-layer corners using U R U' L' U R' U' L'; then permute last-layer edges using  
 Ua/Ub permutations. I will implement helper checks and safe application with inverses so I can revert  
 if a guess doesn't solve. I'll check for solved after each stage and finish with final\_answer("") if  
 solved.

**Step 13****Thought.**

I will correct the cube indexing for edges and corners (earlier mappings had mistakes), then run a full  
 beginner-method pipeline: build the white cross, insert white corners, solve middle-layer edges, and  
 finish with OLL and PLL. I will expand any "2" turns into two quarter turns so the environment accepts  
 them. I will check for solved after each phase, and if solved at any point, I will immediately call  
 final\_answer("") .

**Code.**

```

# Corrected, end-to-end beginner method solver with accurate index mappings and safe "2" expansion.

# ====== Helpers and Mappings ======
def do(m):
    # Expand "2" moves into two quarter turns
    if m.endswith("2"):
        f = m[0]
        make_move(f)
        make_move(f)
    else:
        make_move(m)

def do2(face):
    make_move(face)
    make_move(face)

def is_solved_state(s):
    return (all(c == 'R' for c in s[0:9]) and
            all(c == 'G' for c in s[9:18]) and
            all(c == 'B' for c in s[18:27]) and
            all(c == 'Y' for c in s[27:36]) and
            all(c == 'O' for c in s[36:45]) and
            all(c == 'W' for c in s[45:54]))

def centers_of(state):
    return {
        'F': state[4],    # R
        'B': state[13],   # G
        'L': state[22],   # B
        'R': state[31],   # Y
        'U': state[40],   # O
        'D': state[49],   # W
    }

# Correct edge indices based on the provided net
edge_defs = {

```

```

1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051

    'UF': (('U', 43), ('F', 1)),
    'UR': (('U', 41), ('R', 28)),
    'UB': (('U', 37), ('B', 10)),
    'UL': (('U', 39), ('L', 19)),
    'FR': (('F', 5), ('R', 30)),
    'FL': (('F', 3), ('L', 23)),
    'RB': (('R', 32), ('B', 12)),
    'LB': (('L', 21), ('B', 14)),
    'FD': (('F', 7), ('D', 46)),
    'RD': (('R', 34), ('D', 50)),
    'BD': (('B', 16), ('D', 52)),
    'LD': (('L', 25), ('D', 48)),
}

# Correct corner indices
corners = {
    'UFR': [('U', 44), ('F', 2), ('R', 27)],
    'URB': [('U', 38), ('R', 29), ('B', 9)],
    'UBL': [('U', 36), ('B', 11), ('L', 18)],
    'ULF': [('U', 42), ('L', 20), ('F', 0)],
    'DFR': [('D', 47), ('F', 8), ('R', 33)],
    'DRB': [('D', 45), ('R', 35), ('B', 15)],
    'DBL': [('D', 51), ('B', 17), ('L', 24)],
    'DLF': [('D', 53), ('L', 26), ('F', 6)],
}

# Cycles for U and D layer edges
U_cycle = ['UF', 'UR', 'UB', 'UL']
D_cycle_cw = ['FD', 'RD', 'BD', 'LD'] # Under D (clockwise)

UX_of = {'F': 'UF', 'R': 'UR', 'B': 'UB', 'L': 'UL'}
XD_of = {'F': 'FD', 'R': 'RD', 'B': 'BD', 'L': 'LD'}

def U_times(k):
    for _ in range(k % 4): do('U')

def D_times(k):
    for _ in range(k % 4): do('D')

def align_U_edge_to(state, cur_edge, target_edge):
    ci = U_cycle.index(cur_edge)
    ti = U_cycle.index(target_edge)
    k = (ti - ci) % 4
    U_times(k)

def align_D_edge_to(state, cur_edge, target_edge):
    ci = D_cycle_cw.index(cur_edge)
    ti = D_cycle_cw.index(target_edge)
    k = (ti - ci) % 4
    D_times(k)

def get_edge_colors(state, edge_name):
    (f1, i1), (f2, i2) = edge_defs[edge_name]
    return (state[i1], state[i2])

def locate_white_edge_partner(state, partner_color):
    for name, ((f1,i1), (f2,i2)) in edge_defs.items():
        c1, c2 = state[i1], state[i2]
        if c1 == 'W' and c2 == partner_color:
            return name, 0
        if c2 == 'W' and c1 == partner_color:
            return name, 1
    return None, None

# ===== White Cross =====

def lift_middle_white_to_U(state, edge_name, white_on_face):
    # Move middle edge to U with white on U
    if edge_name == 'FR':
        if white_on_face == 'F': do("F'")
        elif white_on_face == 'R': do("R'")
    elif edge_name == 'FL':
        if white_on_face == 'F': do("F'")
        elif white_on_face == 'L': do("L'")
    elif edge_name == 'RB':
        if white_on_face == 'R': do("R'")
        elif white_on_face == 'B': do("B'")
    elif edge_name == 'LB':
        if white_on_face == 'L': do("L'")
        elif white_on_face == 'B': do("B'")

```

```

2052
2053     def drop_U_side_white_to_D(state, edge_name, white_on_face):
2054         if edge_name == 'UF' and white_on_face == 'F': do("F")
2055         elif edge_name == 'UR' and white_on_face == 'R': do("R")
2056         elif edge_name == 'UB' and white_on_face == 'B': do("B")
2057         elif edge_name == 'UL' and white_on_face == 'L': do("L")
2058
2059     def eject_D_side_white(state, edge_name, white_on_face):
2060         target = white_on_face + 'D'
2061         align_D_edge_to(state, edge_name, target)
2062         do(white_on_face + "")
2063
2064     def insert_from_U_white_up(state, target_face):
2065         desired_u = UX_of[target_face]
2066         st2 = get_observation()
2067         centers = centers_of(st2)
2068         partner = centers[target_face]
2069         current_u = None
2070         for en in U_cycle:
2071             (f1,i1),(f2,i2) = edge_defs[en]
2072             c1, c2 = st2[i1], st2[i2]
2073             if (c1 == 'W' and c2 == partner) or (c2 == 'W' and c1 == partner):
2074                 current_u = en
2075                 break
2076         if current_u is None:
2077             return
2078         align_U_edge_to(st2, current_u, desired_u)
2079         do2(target_face)
2080
2081     def build_white_cross():
2082         order = ['F', 'R', 'B', 'L']
2083         for face in order:
2084             for _ in range(60):
2085                 st = get_observation()
2086                 ctr = centers_of(st)
2087                 partner = ctr[face]
2088                 # Check if solved at XD (white on D index for that edge)
2089                 xd = XD_of[face]
2090                 (f1,i1),(f2,i2) = edge_defs[xd]
2091                 d_idx = i1 if f1 == 'D' else (i2 if f2 == 'D' else None)
2092                 if d_idx is not None:
2093                     if ((st[i1] == partner and st[i2] == 'W') or (st[i2] == partner and st[i1] == 'W')) and
2094                         st[d_idx] == 'W':
2095                             break # Edge in place and oriented
2096                 # Locate target edge
2097                 en, pos = locate_white_edge_partner(st, partner)
2098                 if en is None:
2099                     break
2100                 (ef1,ei1),(ef2,ei2) = edge_defs[en]
2101                 # Determine where white is
2102                 white_on = ef1 if st[ei1] == 'W' else (ef2 if st[ei2] == 'W' else None)
2103                 if en in U_cycle:
2104                     if white_on == 'U':
2105                         insert_from_U_white_up(st, face)
2106                     else:
2107                         drop_U_side_white_to_D(st, en, white_on)
2108                 elif en in ['FR', 'FL', 'RB', 'LB']:
2109                     lift_middle_white_to_U(st, en, white_on)
2110                 elif en in D_cycle_cw:
2111                     if white_on == 'D':
2112                         align_D_edge_to(st, en, XD_of[face])
2113                         do2(face)
2114                     else:
2115                         eject_D_side_white(st, en, white_on)
2116
2117             # ===== White Corners =====
2118
2119             def get_corner_colors(state, name):
2120                 return tuple(state[idx] for _, idx in corners[name])
2121
2122             def corner_color_set(state, name):
2123                 return set(get_corner_colors(state, name))
2124
2125             top_cycle = ['UFR', 'URB', 'UBL', 'ULF']
2126             down_cycle = ['DFR', 'DRB', 'DBL', 'DLF']
2127
2128             def U_to_top_corner(current_top, target_top):
2129                 ci = top_cycle.index(current_top)
2130                 ti = top_cycle.index(target_top)
2131                 U_times((ti - ci) % 4)
2132
2133             def D_to_down_corner(current_down, target_down):
2134                 ci = down_cycle.index(current_down)
2135                 ti = down_cycle.index(target_down)

```

```

2106
2107
2108     D_times((ti - ci) % 4)
2109
2110     def is_corner_solved(state, name):
2111         # White on D face and other two colors match centers
2112         faces = [f for f, _ in corners[name]]
2113         idxs = [i for _, i in corners[name]]
2114         d_pos = None
2115         for i, f in enumerate(faces):
2116             if f == 'D':
2117                 d_pos = i
2118                 break
2119         if d_pos is None:
2120             return False
2121         if state[idxs[d_pos]] != 'W':
2122             return False
2123         ctr = centers_of(state)
2124         for i, f in enumerate(faces):
2125             if f == 'D': continue
2126             if state[idxs[i]] != ctr[f]:
2127                 return False
2128         return True
2129
2130     triggers = {
2131         'DFR': {'top': 'UFR', 'do': lambda: (do('R'), do('U'), do('R')), do('U'))},
2132         'DRB': {'top': 'URB', 'do': lambda: (do('B'), do('U'), do('B')), do('U'))},
2133         'DBL': {'top': 'UBL', 'do': lambda: (do('L'), do('U')), do('L'), do('U'))},
2134         'DLF': {'top': 'ULF', 'do': lambda: (do('F'), do('U')), do('F'), do('U'))},
2135     }
2136
2137     def find_corner_by_colors(state, target_set):
2138         for name in corners:
2139             if corner_color_set(state, name) == target_set:
2140                 return name
2141         return None
2142
2143     def eject_down_corner(slot_name):
2144         trig = triggers[slot_name]['do']
2145         for _ in range(3):
2146             trig()
2147
2148     def insert_corner(slot_name, target_set):
2149         top_name = triggers[slot_name]['top']
2150         trig = triggers[slot_name]['do']
2151         for _ in range(40):
2152             st = get_observation()
2153             if is_corner_solved(st, slot_name):
2154                 return True
2155             pos = find_corner_by_colors(st, target_set)
2156             if pos is None:
2157                 return False
2158             if pos in down_cycle:
2159                 D_to_down_corner(pos, slot_name)
2160                 eject_down_corner(slot_name)
2161             elif pos in top_cycle:
2162                 U_to_top_corner(pos, top_name)
2163                 trig()
2164             else:
2165                 trig()
2166         return False
2167
2168     def solve_white_corners():
2169         st = get_observation()
2170         ctr = centers_of(st)
2171         targets = [
2172             ('DFR', {'W', ctr['F'], ctr['R']}),
2173             ('DRB', {'W', ctr['R'], ctr['B']}),
2174             ('DBL', {'W', ctr['B'], ctr['L']}),
2175             ('DLF', {'W', ctr['L'], ctr['F']}),
2176         ]
2177         for slot, tset in targets:
2178             for _ in range(60):
2179                 st = get_observation()
2180                 if is_corner_solved(st, slot):
2181                     break
2182                 insert_corner(slot, tset)
2183
2184     # ===== Middle Layer =====
2185
2186     def right_of(face):
2187         return {'F': 'R', 'R': 'B', 'B': 'L', 'L': 'F'}[face]
2188
2189     def left_of(face):
2190         return {'F': 'L', 'L': 'B', 'B': 'R', 'R': 'F'}[face]

```

```

2160
2161
2162     def is_middle_edge_solved(state, slot):
2163         (f1, i1), (f2, i2) = edge_defs[slot]
2164         c1, c2 = state[i1], state[i2]
2165         ctr = centers_of(state)
2166         return c1 == ctr[f1] and c2 == ctr[f2]
2167
2168     def any_middle_unsolved(state):
2169         return not (is_middle_edge_solved(state, 'FR') and
2170                     is_middle_edge_solved(state, 'RB') and
2171                     is_middle_edge_solved(state, 'LB') and
2172                     is_middle_edge_solved(state, 'FL'))
2173
2174     def find_u_edge_candidate(state):
2175         for en in U_cycle:
2176             (f1, i1), (f2, i2) = edge_defs[en]
2177             c1, c2 = state[i1], state[i2]
2178             if c1 in ('O', 'W') or c2 in ('O', 'W'):
2179                 continue
2180             return en
2181         return None
2182
2183     def find_u_edge_by_colors(state, colors_set):
2184         for en in U_cycle:
2185             (f1, i1), (f2, i2) = edge_defs[en]
2186             if set([state[i1], state[i2]]) == colors_set:
2187                 return en
2188         return None
2189
2190     def insert_middle_from_U_above(front_face, direction):
2191         if direction == 'right':
2192             rf = right_of(front_face)
2193             do('U'); do(rf); do("U'"); do("U'"); do(front_face+""); do('U'); do(front_face)
2194         else:
2195             lf = left_of(front_face)
2196             do("U'"); do(lf+""); do("U'"); do(lf); do('U'); do(front_face); do("U'"); do(front_face+"")
2197
2198     def eject_incorrect_middle(front_face):
2199         insert_middle_from_U_above(front_face, 'right')
2200
2201     def solve_middle_layer():
2202         for _ in range(300):
2203             st = get_observation()
2204             if not any_middle_unsolved(st):
2205                 break
2206             st = get_observation()
2207             ctr = centers_of(st)
2208             en = find_u_edge_candidate(st)
2209             if en is not None:
2210                 side_map = {'UF': 'F', 'UR': 'R', 'UB': 'B', 'UL': 'L'}
2211                 side = side_map[en]
2212                 (f1, i1), (f2, i2) = edge_defs[en]
2213                 if f1 == side:
2214                     side_color = st[i1]
2215                     up_color = st[i2]
2216                 else:
2217                     side_color = st[i2]
2218                     up_color = st[i1]
2219                 target_front = None
2220                 for f in ['F', 'R', 'B', 'L']:
2221                     if ctr[f] == side_color:
2222                         target_front = f
2223                         break
2224                 if target_front is None:
2225                     do('U'); continue
2226                 st2 = get_observation()
2227                 cur = find_u_edge_by_colors(st2, set([side_color, up_color]))
2228                 if cur is None:
2229                     do('U'); continue
2230                 align_U_edge_to(st2, cur, UX_of[target_front])
2231                 if ctr[right_of(target_front)] == up_color:
2232                     insert_middle_from_U_above(target_front, 'right')
2233                 elif ctr[left_of(target_front)] == up_color:
2234                     insert_middle_from_U_above(target_front, 'left')
2235                 else:
2236                     do('U')
2237             else:
2238                 st = get_observation()
2239                 if not is_middle_edge_solved(st, 'FR'):
2240                     eject_incorrect_middle('F')
2241                 elif not is_middle_edge_solved(st, 'RB'):
2242                     eject_incorrect_middle('R')
2243                 elif not is_middle_edge_solved(st, 'LB'):
2244

```

```

2214
2215         eject_incorrect_middle('L')
2216     elif not is_middle_edge_solved(st, 'FL'):
2217         eject_incorrect_middle('F')
2218
2219     # ===== Last Layer (OLL + PLL) =====
2220
2221     U_indices = [36,37,38,39,40,41,42,43,44]
2222     def u_face_all_orange(state):
2223         return all(state[i] == 'O' for i in U_indices)
2224
2225     def u_edges(state):
2226         return {'top': state[37], 'left': state[39], 'right': state[41], 'bottom': state[43]}
2227
2228     def u_cross_orange(state):
2229         e = u_edges(state)
2230         return e['top']=='O' and e['left']=='O' and e['right']=='O' and e['bottom']=='O'
2231
2232     def solve_oll_edges():
2233         FRURUFp = ["F", "R", "U", "R'", "U'", "F'"]
2234         for _ in range(12):
2235             st = get_observation()
2236             if u_cross_orange(st):
2237                 return
2238             e = u_edges(st)
2239             flags = {pos:(col=='O') for pos,col in e.items()}
2240             cnt = sum(flags.values())
2241             if cnt == 0:
2242                 do_seq(FRURUFp)
2243             elif cnt == 2:
2244                 if flags['left'] and flags['right']:
2245                     do_seq(FRURUFp)
2246                 elif flags['top'] and flags['bottom']:
2247                     U_times(1); do_seq(FRURUFp)
2248                 else:
2249                     for _ in range(4):
2250                         st2 = get_observation()
2251                         e2 = u_edges(st2)
2252                         if e2['top']=='O' and e2['left']=='O':
2253                             do_seq(FRURUFp); break
2254                         U_times(1)
2255                     else:
2256                         do_seq(FRURUFp)
2257
2258     def solve_oll_corners():
2259         sune = ["R'", "U'", "R'", "U", "R", "U2", "R'"]
2260         antisune = ["R'", "U'", "R", "U'", "R'", "U2", "R"]
2261         for _ in range(36):
2262             st = get_observation()
2263             if u_face_all_orange(st):
2264                 return
2265             for _ in range(4):
2266                 st = get_observation()
2267                 if st[2] == 'O':
2268                     do_seq(sune); break
2269                 elif st[2] == 'O':
2270                     do_seq(antisune); break
2271                 else:
2272                     do('U')
2273
2274         edges_U_top = {
2275             'UF': (('U',43), ('F',1)),
2276             'UR': (('U',41), ('R',28)),
2277             'UB': (('U',37), ('B',10)),
2278             'UL': (('U',39), ('L',19)),
2279         }
2280
2281     def edges_in_place_top(state):
2282         ctr = centers_of(state)
2283         for name, ((f1,i1),(f2,i2)) in edges_U_top.items():
2284             if state[i2] != ctr[f2]:
2285                 return False
2286         return True
2287
2288     def corner_side_colors_match(state, corner_name):
2289         faces = [f for f,_ in corners[corner_name]]
2290         idxs = [i for _,i in corners[corner_name]]
2291         ctr = centers_of(state)
2292         for f,i in zip(faces,idxs):
2293             if f == 'U': continue
2294             if state[i] != ctr[f]:
2295                 return False
2296         return True

```

```

2268
2269     def all_corners_in_place(state):
2270         return all(corner_side_colors_match(state, name) for name in ['UFR', 'URB', 'UBL', 'ULF'])
2271
2272     corner_cycle = ["U", "R", "U'", "L'", "U", "R'", "U'", "L"]
2273     Ua = ["R", "U", "R", "U", "R", "U", "R", "U", "R", "U", "R", "U"]
2274     Ub = ["R2", "U", "R", "U", "R'", "U'", "R'", "U", "R", "U", "R"]
2275
2276     def solve_pll_corners():
2277         for _ in range(12):
2278             st = get_observation()
2279             if all_corners_in_place(st):
2280                 return
2281             # try to position a correct corner at UFR
2282             for _ in range(4):
2283                 st = get_observation()
2284                 if corner_side_colors_match(st, 'UFR'):
2285                     break
2286                 do('U')
2287                 do_seq(corner_cycle)
2288
2289     def invert_move(m):
2290         if m.endswith("2"): return m
2291         if m.endswith(")": return m[:-1]
2292         return m + ")"
2293
2294     def invert_seq(seq):
2295         return [invert_move(m) for m in reversed(seq)]
2296
2297     def try_edge_perm(alg):
2298         do_seq(alg)
2299         st = get_observation()
2300         if is_solved_state(st) or edges_in_place_top(st):
2301             return True
2302         do_seq(invert_seq(alg))
2303         return False
2304
2305     def solve_pll_edges():
2306         for _ in range(12):
2307             st = get_observation()
2308             if edges_in_place_top(st):
2309                 return
2310             for _ in range(4):
2311                 if try_edge_perm(Ua): return
2312                 if try_edge_perm(Ub): return
2313                 do('U')
2314
2315     # ===== Run the full pipeline =====
2316
2317     st0 = get_observation()
2318     if is_solved_state(st0):
2319         final_answer("")
2320     else:
2321         # 1) White cross
2322         build_white_cross()
2323         if is_solved_state(get_observation()):
2324             final_answer("")
2325         # 2) White corners
2326         solve_white_corners()
2327         if is_solved_state(get_observation()):
2328             final_answer("")
2329         # 3) Middle layer
2330         solve_middle_layer()
2331         if is_solved_state(get_observation()):
2332             final_answer("")
2333         # 4) Last layer
2334         solve_oll_edges()
2335         solve_oll_corners()
2336         solve_pll_corners()
2337         solve_pll_edges()
2338         # Align U and check
2339         for _ in range(4):
2340             if is_solved_state(get_observation()):
2341                 final_answer("")
2342             do('U')
2343         # If still not solved, print state for further iteration
2344         print(get_observation())
2345
2346     </code>

```

## Observation.

2322  
 2323  
 2324 RWRORRGGRGRGGWGORYBYBGBOBYGGOYYYYYBBOOWOBOORWBWWWW  
 2325  
 2326  
 2327

### 2328 L.3 BASIC AGENT – BRUTE-FORCE ENUMERATION

2329	Model	Agent Type	Input	Reward	Depth
2330	gpt-oss-120b	Basic	Full Symbolic	No reward	12

2332  
 2333 This case demonstrates the behavior of a less capable agent in a no-reward setting. The agent executes up to 20,000 moves, selecting each move essentially at random, with only a minimal stochastic backtracking mechanism (5% chance of undoing the previous move). There is no systematic planning, pruning, or search applied, and the cube remains unsolved. Aside from the stochastic backtracking, the agent’s behavior is equivalent to a random walk over the state space.  
 2337

#### 2338 Code.

```
2339
2340 import random
2341
2342 # Moves and their inverses
2343 moves = ["F", "F'", "B", "B'", "L", "L'", "R", "R'", "U", "U'", "D", "D'"]
2344 inverse = {"F": "F'", "F'": "F", "B": "B'", "B'": "B",
2345     "L": "L'", "L'": "L", "R": "R'", "R'": "R",
2346     "U": "U'", "U'": "U", "D": "D'", "D'": "D"}
2347
2348 def is_solved(state):
2349     # solved state order: FRONT=R, BACK=G, LEFT=B, RIGHT=Y, UP=O, DOWN=W
2350     solved = "R"*9 + "G"*9 + "B"*9 + "Y"*9 + "O"*9 + "W"*9
2351     return state == solved
2352
2353 # Simulated annealing / random walk with occasional acceptance of worse moves
2354 max_steps = 20000 # allow many moves within a single tool call
2355 state = get_observation()
2356 if is_solved(state):
2357     print("Already solved")
2358 else:
2359     for step in range(max_steps):
2360         if is_solved(state):
2361             print("Cube solved at step", step)
2362             break
2363         mv = random.choice(moves)
2364         make_move(mv)
2365         new_state = get_observation()
2366         # If the move improved (or equal) heuristic, keep it.
2367         # We don't have direct heuristic, so we just keep the move most of the time,
2368         # but occasionally revert a bad move to avoid dead-ends.
2369         # Simple rule: if new_state is same as previous (unlikely) revert, else keep.
2370         # To allow some backtracking, we revert with small probability.
2371         keep = True
2372         if random.random() < 0.05: # 5% chance to undo
2373             make_move(inverse[mv])
2374             new_state = get_observation()
2375             keep = False
2376         state = new_state
2377     else:
2378         print("Reached max steps without solving. Final state:", state)
2379
2380 # Final verification
2381 final_answer("")
```

2382 Not solved.  
 2383

### 2384 L.4 BASIC AGENT – FAILED PRUNING

2385	Model	Agent Type	Input	Reward	Depth
2386	GPT-5	Basic	Full Symbolic	No reward	12

2376 In this test case, the model implements a *meet-in-the-middle* search for Rubik's Cube solving. The  
 2377 procedure first builds a dictionary of states reachable from the solved cube (`build_goal_map`),  
 2378 and then performs a forward search from the scrambled cube (`forward_meet_in_middle`) to  
 2379 detect an intersection. This reflects the model's recall of classical bidirectional search techniques.  
 2380

2381 The implementation introduces several forms of pruning:

2382 **Move-level pruning.** Immediate inverses are skipped. (`if m == inverse_move(last_move): continue`).  
 2383 These rules eliminate trivial local redundancies such as  $R R'$ , or  $U U'$ .

2384 **State duplicate elimination.** Both functions maintain a `visited` set  
 2385 (`if new_state in visited: continue`) to prevent rediscovering identical states along different paths.  
 2386

2387 Despite these measures, the pruning remains insufficient. Move-level constraints only remove redundant short patterns, and the `visited` set curbs local repetition but cannot control the exponential  
 2388 growth of unique cube states. Indeed, even at depth 5 the algorithm reports over 9,000 states in the  
 2389 goal-side map. Without stronger heuristics (e.g., pattern databases or domain-specific evaluation  
 2390 functions), the search rapidly exhausts computational limits without finding a solution.  
 2391

2392 This case illustrates that while the model can recall canonical pruning strategies, their limited scope  
 2393 fails to achieve practical efficiency in large combinatorial spaces.

### 2394 Thought.

2395  
 2396 The previous searches exceeded operation limits. I will switch to a *meet-in-the-middle* search:  
 2397 - Build a dictionary of states reachable from the solved state within depth  $Dg=5$  (using the learned  
 2398 move permutations), storing the sequences from solved to each state.  
 2399 - Then breadth-first search from the current state up to depth  $Ds=5$ , checking at each node if it  
 2400 intersects the goal-side dictionary. If so, concatenate the forward path with the inverse of the  
 2401 goal-side path to obtain a full solution.  
 2402 - Apply the resulting sequence to the real cube using `make_move`, verify if solved, and if so call  
 2403 `final_answer("")`.  
 2404  
 2405 This approach keeps expansions manageable and avoids heavy recursion or long-running loops.

### 2406 Code.

```
2407
2408 import time
2409
2410 # We assume these exist from earlier:
2411 # perm_by_move, apply_perm, inverse_move, invert_seq, solved, do_moves, get_observation
2412
2413 move_order = ['U', 'U''', 'R', 'R''', 'F', 'F''', 'L', 'L''', 'D', 'D''', 'B', 'B'''']
2414
2415 def build_goal_map(depth_limit=5, time_limit=6.0):
2416     start_time = time.time()
2417     # Each node: (state, last_move, seq_tuple)
2418     start = solved
2419     visited = set([start])
2420     goal_map = {} # state -> seq from solved to state (tuple of moves)
2421     frontier = [(start, None, ())]
2422     depth = 0
2423     while depth < depth_limit and frontier:
2424         if time.time() - start_time > time_limit:
2425             break
2426         next_frontier = []
2427         for state, last_move, seq in frontier:
2428             for m in move_order:
2429                 if last_move is not None:
2430                     # prune: same face twice or immediate inverse
2431                     if m[0] == last_move[0] or m == inverse_move(last_move):
2432                         continue
2433                     new_state = apply_perm(state, perm_by_move[m])
2434                     if new_state in visited:
2435                         continue
2436                     visited.add(new_state)
2437                     new_seq = seq + (m,)
2438                     goal_map[new_state] = new_seq
2439                     next_frontier.append((new_state, m, new_seq))
2440         frontier = next_frontier
2441         depth += 1
```

```

2430
2431     print("Goal-map size (<= depth", depth_limit, "):", len(goal_map))
2432     return goal_map
2433
2434     def forward_meet_in_middle(start_state, goal_map, depth_limit=5, time_limit=6.0):
2435         start_time = time.time()
2436         if start_state in goal_map:
2437             # Directly within goal radius
2438             tail = invert_seq(list(goal_map[start_state]))
2439             return True, tail
2440
2441         visited = set([start_state])
2442         frontier = [(start_state, None, ())]
2443         depth = 0
2444         while depth < depth_limit and frontier:
2445             if time.time() - start_time > time_limit:
2446                 break
2447             next_frontier = []
2448             for state, last_move, seq in frontier:
2449                 if state in goal_map:
2450                     tail = invert_seq(list(goal_map[state]))
2451                     return True, list(seq) + tail
2452                 for m in move_order:
2453                     if last_move is not None:
2454                         if m[0] == last_move[0] or m == inverse_move(last_move):
2455                             continue
2456                         new_state = apply_perm(state, perm_by_move[m])
2457                         if new_state in visited:
2458                             continue
2459                         visited.add(new_state)
2460                         new_seq = seq + (m,)
2461                         # Early meet check
2462                         if new_state in goal_map:
2463                             tail = invert_seq(list(goal_map[new_state]))
2464                             return True, list(new_seq) + tail
2465                         next_frontier.append((new_state, m, new_seq))
2466             frontier = next_frontier
2467             depth += 1
2468         return False, None
2469
2470     # Build meet-in-the-middle structures
2471     start_state = get_observation()
2472     print("Current state:", start_state)
2473     goal_map = build_goal_map(depth_limit=5, time_limit=5.0)
2474     ok, solution_seq = forward_meet_in_middle(start_state, goal_map, depth_limit=5, time_limit=5.0)
2475
2476     if not ok or solution_seq is None:
2477         print("First meet-in-the-middle attempt failed; trying slightly deeper on forward side...")
2478         ok, solution_seq = forward_meet_in_middle(start_state, goal_map, depth_limit=6, time_limit=6.0)
2479
2480     if not ok or solution_seq is None:
2481         print("Could not find a meet-in-the-middle solution within limits.")
2482         final_answer("Could not solve within limits.")
2483     else:
2484         print("Found sequence length:", len(solution_seq))
2485         print("Sequence:", ''.join(solution_seq))
2486         # Apply to real cube
2487         do_moves(solution_seq)
2488         final_state = get_observation()
2489         print("Solved?", final_state == solved)
2490         if final_state == solved:
2491             final_answer("")
2492         else:
2493             final_answer("Applied a sequence, but cube not solved.")
2494
2495

```

## Observation.

2478 Current state: RGGRRWRGWOWBGGBGBOYRRYYBRYYBRGYGBBGWOYOOOWOOROWRBBW  
2479 Goal-map size (<= depth 5 ): 9732  
2480 First meet-in-the-middle attempt failed; trying slightly deeper on forward side...  
<Killed by the Python Interpreter>

*Not solved.*

2484  
2485

## L.5 REWARD – BYPASS VISUAL REASONING

2486  
2487  
2488

Model	Agent Type	Input	Reward	Depth
GPT-5	Basic	Vertex View	Heuristic	3

2489  
2490  
2491  
2492  
2493  
2494

In this case (Step 1), the agent was required to explicitly call `get_observation()` to view the cube image. However, throughout its search process the agent did not make any use of the visual input; it only called the image observation once at the end, to display the final cube state. Instead, its entire search relied exclusively on the heuristic reward. The code shows a greedy lookahead strategy up to depth three, where move sequences were evaluated and selected solely on the basis of their reward values.

2495  
2496  
2497  
2498  
2499

This behavior illustrates that when a simple external reward is available, even a capable agent may effectively bypass the provided visual input. Rather than engaging in visual reasoning to reconstruct the cube state, the agent defaulted to symbolic search in the move space, guided entirely by reward optimization.

2500  
2501  
2502

*Step 1*

**Thought.**

2503  
2504  
2505  
2506  
2507  
2508  
2509  
2510  
2511  
2512  
2513  
2514  
2515

```
Thought: I'll use greedy and limited lookahead search guided by the heuristic reward. At each iteration
I will:
- Try all 12 basic moves and pick the one with the largest positive reward; if found, apply it and
continue.
- If no single move improves, try pairs of moves (depth 2) to find a positive net improvement; if
found, apply the best pair.
- If still no improvement, try depth 3 sequences for a positive net improvement; if found, apply the
best triple.
- Repeat until no improvements are found up to depth 3 or a reasonable move limit is reached.
I will then show the current cube observation so I can visually assess if it's solved. If it is, I will
call final_answer("") .
```

**Code.**

2516  
2517  
2518  
2519  
2520  
2521  
2522  
2523  
2524  
2525  
2526  
2527  
2528  
2529  
2530  
2531  
2532  
2533  
2534  
2535  
2536  
2537

```
# Greedy + lookahead solver using the heuristic reward

# Define moves and inverses
moves = ["F", "F'", "B", "B'", "L", "L'", "R", "R'", "U", "U'", "D", "D'"]
inv = {m: (m[:-1] if m.endswith(')') else m + ')') for m in moves}

def try_best_sequence(max_depth=3):
    """
    Try to find the best improving sequence up to max_depth.
    Returns (best_reward, best_sequence) with best_reward > 0 if found, else (0, []).
    """
    best_reward = 0.0
    best_seq = []

    # Depth 1
    for m1 in moves:
        r1 = make_move(m1)
        # undo
        _ = make_move(inv[m1])
        if r1 > best_reward:
            best_reward = r1
            best_seq = [m1]

    if best_reward > 0 or max_depth == 1:
        return best_reward, best_seq

    # Depth 2
    for m1 in moves:
        r1 = make_move(m1)
        for m2 in moves:
            r2 = make_move(m2)
            net = r1 + r2
            # undo m2
            _ = make_move(inv[m2])
            if net > best_reward:
                best_reward = net
                best_seq = [m1, m2]

    if best_reward > 0 or max_depth == 2:
        return best_reward, best_seq

    # Depth 3
    for m1 in moves:
        r1 = make_move(m1)
        for m2 in moves:
            r2 = make_move(m2)
            for m3 in moves:
                r3 = make_move(m3)
                net = r1 + r2 + r3
                # undo m3
                _ = make_move(inv[m3])
                if net > best_reward:
                    best_reward = net
                    best_seq = [m1, m2, m3]

    if best_reward > 0 or max_depth == 3:
        return best_reward, best_seq
```

```

2538
2539         if net > best_reward:
2540             best_reward = net
2541             best_seq = [m1, m2]
2542             # undo m1
2543             _ = make_move(inv[m1])
2544
2545             if best_reward > 0 or max_depth == 2:
2546                 return best_reward, best_seq
2547
2548             # Depth 3
2549             for m1 in moves:
2550                 r1 = make_move(m1)
2551                 for m2 in moves:
2552                     r2 = make_move(m2)
2553                     for m3 in moves:
2554                         r3 = make_move(m3)
2555                         net = r1 + r2 + r3
2556                         # undo m3
2557                         _ = make_move(inv[m3])
2558                         if net > best_reward:
2559                             best_reward = net
2560                             best_seq = [m1, m2, m3]
2561                         # undo m2
2562                         _ = make_move(inv[m2])
2563                         # undo m1
2564                         _ = make_move(inv[m1])
2565
2566             return best_reward, best_seq
2567
2568             # Main improvement loop
2569             total_improvement = 0.0
2570             applied_moves = []
2571             iterations = 0
2572             max_iterations = 200 # safety cap on number of applied moves
2573
2574             print("Starting heuristic-guided improvement...")
2575             while iterations < max_iterations:
2576                 iterations += 1
2577                 best_r, best_seq = try_best_sequence(max_depth=3)
2578                 if best_r <= 0 or not best_seq:
2579                     print(f"No improving sequence found at iteration {iterations}. Stopping local search.")
2580                     break
2581
2582                 # Apply the best found sequence
2583                 seq_r_sum = 0.0
2584                 for m in best_seq:
2585                     r = make_move(m)
2586                     seq_r_sum += r
2587                     applied_moves.append(m)
2588                 total_improvement += seq_r_sum
2589                 print(f"Iteration {iterations}: Applied {best_seq} with net reward {seq_r_sum:.3f}. Total improvement: {total_improvement:.3f}")
2590
2591             print(f"Finished. Applied {len(applied_moves)} moves. Total heuristic improvement from start: {total_improvement:.3f}")
2592             print("Final observation for visual check:")
2593             obs = get_observation()
2594             print(obs) # This will display the current cube in vertex view for visual verification.
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642
2643
2644
2645
2646
2647
2648
2649
2650
2651
2652
2653
2654
2655
2656
2657
2658
2659
2660
2661
2662
2663
2664
2665
2666
2667
2668
2669
2670
2671
2672
2673
2674
2675
2676
2677
2678
2679
2680
2681
2682
2683
2684
2685
2686
2687
2688
2689
2690
2691

```

## Observation.

```

Starting heuristic-guided improvement...
Iteration 1: Applied ['D', "B'"] with net reward 4.000. Total improvement: 4.000
Iteration 2: Applied ["D'"] with net reward 2.000. Total improvement: 6.000
No improving sequence found at iteration 3. Stopping local search.
Finished. Applied 3 moves. Total heuristic improvement from start: 6.000
Final observation for visual check:
<PIL.Image.Image image mode=RGB size=84x84 at 0x75CD29633A50>

```

## L.6 REWARD – AS A MISLEADING SIGNAL

Model	Agent Type	Input	Reward	Depth
GPT-5	Basic	Full Symbolic	Heuristic	3

2592 This experiment shows how a simplistic external reward can misalign with an advanced agent's  
 2593 internal reasoning. In Steps 2–7, GPT-5 applied several non-trivial search strategies: greedy hill-  
 2594 climbing with depth-2 lookahead and random shaking, heuristic-guided DFS, constrained stochastic  
 2595 search, and macro-based methods incorporating standard Rubik's Cube algorithms (e.g., "Sexy  
 2596 Move" and "Sune"). However, all of these approaches were centered on maximizing the provided  
 2597 heuristic reward—a single numerical score intended to measure progress. The result was limited:  
 2598 the heuristic plateaued (no improvement in `cumulative_delta`), and the cube was never solved.

2599 In Step 8, the agent explicitly reflected, "*The heuristic-based approaches failed to make progress... I*  
 2600 *will switch to a deterministic, constraint-driven search that directly targets subgoals (e.g., build the*  
 2601 *white cross).*" After abandoning the external reward and instead adopting a subgoal-based strategy,  
 2602 the agent made immediate progress (white cross success).

2603 This case illustrates that even for highly capable agents, a simplistic reward signal may conflict  
 2604 with their internal problem-solving process, reducing performance until the agent reverts to its own  
 2605 structured reasoning.

2606 Due to space constraints, we only include the observation outputs and code fragments of the key  
 2607 steps.

2609 **Step 2**

2610 **Thought.**

2612  
 2613 Thought: I will try to solve the cube using the available tools. Since we can get a heuristic reward  
 2614 for each move, I'll implement a greedy hill-climbing approach guided by this reward: at each step,  
 2615 evaluate all moves and choose the one that gives the highest positive improvement. If no single move  
 2616 gives improvement, I'll try a depth-2 lookahead (pairs of moves). If still stuck, I'll use occasional  
 2617 random \shake" moves to escape local minima. I will keep checking if the cube is solved by verifying  
 2618 each face is uniform; once solved, I will call `final_answer("")` as instructed.

2619 **Code.**

```
2621 from random import choice, randint, seed
2622 seed(42)

2623 # Helper functions
2624 def is_solved(state):
2625     if not isinstance(state, str) or len(state) != 54:
2626         return False
2627     faces = [
2628         state[0:9],    # FRONT (R)
2629         state[9:18],   # BACK (G)
2630         state[18:27],  # LEFT (B)
2631         state[27:36],  # RIGHT (Y)
2632         state[36:45],  # UP (O)
2633         state[45:54],  # DOWN (W)
2634     ]
2635     return all(len(set(face)) == 1 for face in faces)

2636 def inverse(move):
2637     return move[:-1] if move.endswith("')") else move + "'

2638 moves = ['F', "F'", 'B', "B'", 'L', "L'", 'R', "R'", 'U', "U'", 'D', "D'"]

2639 def try_single_moves():
2640     # Evaluate all single moves and return the best positive move and its reward
2641     best_move = None
2642     best_reward = float('-inf')
2643     for m in moves:
2644         r = make_move(m)
2645         # revert
2646         make_move(inverse(m))
2647         if r > best_reward:
2648             best_reward = r
2649             best_move = m
2650     return best_move, best_reward

2651 def try_depth2_moves():
2652     # Evaluate all pairs of moves and return the best positive pair and cumulative reward
2653     best_pair = None
```

```

2646
2647     best_reward = float('-inf')
2648     for m1 in moves:
2649         r1 = make_move(m1)
2650         for m2 in moves:
2651             r2 = make_move(m2)
2652             total = r1 + r2
2653             # revert m2 then m1
2654             make_move(inverse(m2))
2655             if total > best_reward:
2656                 best_reward = total
2657                 best_pair = (m1, m2)
2658             make_move(inverse(m1))
2659     return best_pair, best_reward
2660
2661     def faces_solved_count(state):
2662         # number of fully solved faces
2663         faces = [
2664             state[0:9], state[9:18], state[18:27], state[27:36], state[36:45], state[45:54]
2665         ]
2666         return sum(1 for f in faces if len(set(f)) == 1)
2667
2668     # Start from current observation
2669     state0 = get_observation()
2670     print("Initial state:", state0)
2671     print("Initially solved?", is_solved(state0))
2672     print("Initial faces solved:", faces_solved_count(state0))
2673
2674     max_iterations = 600
2675     stagnation = 0
2676     moves_made = []
2677     cumulative_delta = 0.0
2678
2679     for it in range(1, max_iterations + 1):
2680         # Check if solved
2681         st = get_observation()
2682         if is_solved(st):
2683             print(f"Solved at iteration {it} with {len(moves_made)} moves.")
2684             final_answer("") # Set the final answer
2685             break
2686
2687         # Try best single move
2688         m1, r1 = try_single_moves()
2689         if r1 > 0:
2690             rr = make_move(m1)
2691             cumulative_delta += rr
2692             moves_made.append(m1)
2693             stagnation = 0
2694         else:
2695             # Try depth-2
2696             pair, rpair = try_depth2_moves()
2697             if pair is not None and rpair > 0:
2698                 rr1 = make_move(pair[0])
2699                 rr2 = make_move(pair[1])
2700                 cumulative_delta += (rr1 + rr2)
2701                 moves_made.extend([pair[0], pair[1]])
2702                 stagnation = 0
2703             else:
2704                 # Random shake to escape local minima
2705                 rand_moves = randint(1, 3)
2706                 for _ in range(rand_moves):
2707                     m = choice(moves)
2708                     rr = make_move(m)
2709                     cumulative_delta += rr
2710                     moves_made.append(m)
2711                 stagnation += 1
2712
2713         # Periodic logging
2714         if it % 20 == 0:
2715             st = get_observation()
2716             print(f"Iter {it}: cumulative_delta={cumulative_delta:.2f}, moves={len(moves_made)}, faces_solved={faces_solved_count(st)}")
2717
2718     else:
2719         # If loop completes without break (not solved), print status
2720         st = get_observation()
2721         print("Not solved within iteration budget.")
2722         print("Faces solved:", faces_solved_count(st))
2723         print("Current state:", st)

```

## Observation.

```

2700
2701
2702     Initial state: BRRBRRWRRYYOGGGGGGBBWBWGGYBYYBYYBRROOOOOGWWRROOWWWYWWY
2703     Initially solved? False
2704     Initial faces solved: 0
2705     Iter 20: cumulative_delta=0.00, moves=33, faces_solved=0
2706     Iter 40: cumulative_delta=0.00, moves=71, faces_solved=0
2707     Iter 60: cumulative_delta=1.00, moves=105, faces_solved=0
2708     Iter 80: cumulative_delta=0.00, moves=147, faces_solved=0
2709     Iter 100: cumulative_delta=0.00, moves=189, faces_solved=0
2710     Iter 120: cumulative_delta=0.00, moves=230, faces_solved=0
2711     Iter 140: cumulative_delta=0.00, moves=267, faces_solved=0
2712     Iter 160: cumulative_delta=0.00, moves=312, faces_solved=0
2713     Iter 180: cumulative_delta=0.00, moves=351, faces_solved=0
2714     Iter 200: cumulative_delta=0.00, moves=387, faces_solved=0
2715     Iter 220: cumulative_delta=0.00, moves=431, faces_solved=0
2716     Iter 240: cumulative_delta=0.00, moves=472, faces_solved=0
2717     Iter 260: cumulative_delta=0.00, moves=512, faces_solved=0
2718     Iter 280: cumulative_delta=0.00, moves=555, faces_solved=0
2719     Iter 300: cumulative_delta=0.00, moves=595, faces_solved=0
2720     Iter 320: cumulative_delta=0.00, moves=638, faces_solved=0
2721     Iter 340: cumulative_delta=0.00, moves=678, faces_solved=0
2722     Iter 360: cumulative_delta=0.00, moves=718, faces_solved=0
2723     Iter 380: cumulative_delta=0.00, moves=757, faces_solved=0
2724     Iter 400: cumulative_delta=0.00, moves=796, faces_solved=0
2725     Iter 420: cumulative_delta=0.00, moves=834, faces_solved=0
2726     Iter 440: cumulative_delta=0.00, moves=877, faces_solved=0
2727     Iter 460: cumulative_delta=0.00, moves=922, faces_solved=0
2728     Iter 480: cumulative_delta=0.00, moves=965, faces_solved=0
2729     Iter 500: cumulative_delta=0.00, moves=1000, faces_solved=0
2730     Iter 520: cumulative_delta=0.00, moves=1043, faces_solved=0
2731     Iter 540: cumulative_delta=0.00, moves=1082, faces_solved=0
2732     Iter 560: cumulative_delta=0.00, moves=1124, faces_solved=0
2733     Iter 580: cumulative_delta=0.00, moves=1167, faces_solved=0
2734     Iter 600: cumulative_delta=0.00, moves=1210, faces_solved=0

```

### Step 3

#### Thought.

The previous greedy random approach did not make reliable progress in the heuristic metric. I will switch to a monotonic heuristic-guided depth-limited search: at each iteration, I will search for a short sequence of moves (up to a given depth) whose cumulative heuristic reward is non-negative and preferably positive. This avoids regressing in the staged heuristic and should help climb from 0 → 7. I will implement a DFS that only accepts moves with non-negative immediate reward, avoids repeating the same face consecutively, and finds the best cumulative gain sequence. I will apply sequences that yield positive gain; if none found, I can apply a neutral (0 gain) sequence to reposition and try again with deeper depth. I will keep checking for solved state, and call `final_answer("")` immediately once solved.

### Step 4

#### Thought.

The previous approach explored too many branches and hit the operation limit. I'll switch to a constrained stochastic search that samples a limited number of short random sequences each iteration, evaluates their net heuristic gain (sum of move rewards), and applies the best positive-gain sequence if found. If none is found, I'll apply a small random "shake" to escape local minima. This keeps computation bounded and avoids deep recursion. I'll regularly check if the cube is solved and call `final_answer("")` immediately when it is.

### Step 5

#### Thought.

2754

2755

2756 The previous stochastic and exhaustive searches were too heavy and hit the operation limit without  
 2757 improving the heuristic. I will switch to a lightweight, structured macro-search: try a small library  
 2758 of well-known Rubik's Cube algorithms (sexy move, sledgehammer, sune, etc.). For each macro, I'll  
 2759 evaluate its net heuristic gain by executing it (summing rewards). If net gain > 0, I'll keep it;  
 2760 otherwise I'll undo it. This bounded, deterministic approach reduces operations while still seeking  
 2761 monotonic heuristic improvements. I'll repeat this for a limited number of iterations, checking for  
 2762 solved state frequently. If solved, I will immediately call `final_answer("")`.

2763

2764

2765

### Step 8

#### Thought.

2766

2767

2768

2769

2770

2771

2772

2773

2774

2775

The heuristic-based approaches failed to make progress and hit operation limits. I will switch to a deterministic, constraint-driven search that directly targets subgoals (e.g., build the white cross on the Down face) by exploring short sequences using the environment as the cube model. The search:  
 - Applies moves via `make_move` and backtracks with inverse moves to keep the environment consistent.  
 - Prunes immediate inverses and same-face repetitions.  
 - Checks goal predicates by reading the `state_string` via `get_observation()`.  
 I'll first try to complete the white cross (four Down edges in correct positions and orientations). If I get the cube solved at any point, I'll immediately call `final_answer("")`.

#### Code.

2776

2777

2778

2779

2780

2781

2782

2783

2784

2785

2786

2787

2788

2789

2790

2791

2792

2793

2794

2795

2796

2797

2798

2799

2800

2801

2802

2803

2804

2805

2806

2807

```
import random

# Utilities
def is_solved(state):
    if not isinstance(state, str) or len(state) != 54:
        return False
    faces = [state[0:9], state[9:18], state[18:27], state[27:36], state[36:45], state[45:54]]
    return all(len(set(face)) == 1 for face in faces)

def faces_solved_count(state):
    faces = [state[0:9], state[9:18], state[18:27], state[27:36], state[36:45], state[45:54]]
    return sum(1 for f in faces if len(set(f)) == 1)

def inverse_move(m):
    return m[:-1] if m.endswith('') else m + ""

moves = ['F', 'F\'', 'B', 'B\'', 'L', 'L\'', 'R', 'R\'', 'U', 'U\'', 'D', 'D\'']

# Target constraints helpers: (index -> expected color)
# Face centers for reference (fixed)
CENTER = { 'F': 4, 'B': 13, 'L': 22, 'R': 31, 'U': 40, 'D': 49 }
CENTER_COLOR = { 4:'R', 13:'G', 22:'B', 31:'Y', 40:'O', 49:'W' }

# For the white cross on Down (D), we require the four D edges to be placed and oriented correctly.
# Based on the provided net:
# Down face indices:
# 45 46 47
# 48 49 50
# 51 52 53
# Adjacent matching edges on side faces:
# DF: D[0][1]=46 with F[2][1]=7 -> colors W and R
# DR: D[1][2]=50 with R[2][1]=34 -> colors W and Y
# DB: D[2][1]=52 with B[0][1]=10 -> colors W and G
# DL: D[1][0]=48 with L[2][1]=25 -> colors W and B

def cross_constraints():
    return {
        46: 'W', 7: 'R',
        50: 'W', 34: 'Y',
        52: 'W', 10: 'G',
        48: 'W', 25: 'B',
    }

def edge_constraint_DF():
    return {46:'W', 7:'R'}

def edge_constraint_DR():
    return {50:'W', 34:'Y'}
```

```

2808
2809
2810     def edge_constraint_DB():
2811         return {52:'W', 10:'G'}
2812
2813     def edge_constraint_DL():
2814         return {48:'W', 25:'B'}
2815
2816     def constraints_satisfied(state, constraints):
2817         # constraints: dict index->color
2818         for idx, col in constraints.items():
2819             if state[idx] != col:
2820                 return False
2821         return True
2822
2823     def get_state():
2824         return get_observation()
2825
2826     # Depth-limited DFS to find a sequence that satisfies a given constraints set at the end.
2827     # Prunes immediate inverse and same-face repetitions.
2828     def find_sequence_for_constraints(constraints, max_depth=7, max_expansions=120000):
2829         best = {'seq': None}
2830         counters = {'exp': 0}
2831         path = []
2832
2833         def dfs(depth, last_move):
2834             if best['seq'] is not None:
2835                 return
2836             if depth > max_depth or counters['exp'] >= max_expansions:
2837                 return
2838             # Check immediate goal (depth can be 0 at start)
2839             st = get_state()
2840             if constraints_satisfied(st, constraints):
2841                 best['seq'] = path.copy()
2842                 return
2843             if depth == max_depth:
2844                 return
2845             for m in moves:
2846                 if last_move is not None:
2847                     if m[0] == last_move[0]:      # avoid same face twice in a row
2848                         continue
2849                     if m == inverse_move(last_move): # avoid immediate inverse
2850                         continue
2851                     r = make_move(m)
2852                     counters['exp'] += 1
2853                     path.append(m)
2854                     dfs(depth + 1, m)
2855                     if best['seq'] is not None:
2856                         # leave state as is; sequence already applied during search? No, we applied the moves
2857                         # during exploration.
2858                         # We must backtrack to restore before returning control to caller
2859                         make_move(inverse_move(m))
2860                         path.pop()
2861                     return
2862             # backtrack this move
2863             make_move(inverse_move(m))
2864             path.pop()
2865             if counters['exp'] >= max_expansions or best['seq'] is not None:
2866                 return
2867
2868         dfs(0, None)
2869         return best['seq']
2870
2871     # Function to attempt building the white cross iteratively
2872     def build_white_cross():
2873         targets = [edge_constraint_DF(), edge_constraint_DR(), edge_constraint_DB(), edge_constraint_DL()]
2874         # We'll try each edge until all are solved
2875         attempts_per_edge = 6 # bounded attempts with small random perturbations
2876         for i, tgt in enumerate(targets):
2877             for attempt in range(attempts_per_edge):
2878                 st = get_state()
2879                 # If edge already solved, move to next
2880                 if constraints_satisfied(st, tgt):
2881                     break
2882                 # Find a short sequence to solve this edge
2883                 seq = find_sequence_for_constraints(tgt, max_depth=6, max_expansions=90000)
2884                 if seq is not None:
2885                     # Apply the sequence (it's already applied during the search? No, we backtracked; we
2886                     # need to apply now)
2887                     for m in seq:
2888                         make_move(m)
2889                     # Verify
2890                     st2 = get_state()
2891                     if not constraints_satisfied(st2, tgt):
2892

```

```

2862
2863         # If somehow not satisfied, undo and try a tiny shake
2864         for m in reversed(seq):
2865             make_move(inverse_move(m))
2866         else:
2867             break
2868         else:
2869             # Small random shake to change positions
2870             k = random.randint(1, 2)
2871             last_face = None
2872             for _ in range(k):
2873                 cand = [m for m in moves if last_face is None or m[0] != last_face]
2874                 mv = random.choice(cand)
2875                 make_move(mv)
2876                 last_face = mv[0]
2877             else:
2878                 # Failed to solve this edge within attempts
2879             return False
2880
2881         # After solving all four individual edges, ensure full cross correctness (all 8 constraints)
2882         st = get_state()
2883         if constraints_satisfied(st, cross_constraints()):
2884             return True
2885
2886         # If individual edges are placed but cross not fully consistent, try a short refinement to meet
2887         # full constraints
2888         seq = find_sequence_for_constraints(cross_constraints(), max_depth=7, max_expansions=140000)
2889         if seq is not None:
2890             for m in seq:
2891                 make_move(m)
2892             st2 = get_state()
2893             return constraints_satisfied(st2, cross_constraints())
2894         return False
2895
2896         # Driver: attempt to build white cross; if cube becomes solved, finalize.
2897         random.seed(23)
2898         state0 = get_state()
2899         print("Start state:", state0)
2900         print("Faces solved at start:", faces_solved_count(state0))
2901         if is_solved(state0):
2902             final_answer("")
2903
2904         success_cross = build_white_cross()
2905         st = get_state()
2906         print("White cross success:", success_cross, "Faces solved after cross attempt:",
2907               faces_solved_count(st), "State:", st)
2908
2909         # If by chance the cube is fully solved now, finalize.
2910         if is_solved(st):
2911             final_answer("")
2912
2913
2914
2915

```

## Observation.

```

Start state: GYOGROWRBOWWGGBORRROOBBRGYYRWBYRWBGYBGOOBGRWGYWWBYY
Faces solved at start: 0
White cross success: True Faces solved after cross attempt: 0 State:
GYRGRRRRWGGWYGORBBYROGBBYBOWWYYBYYRORGOOGBOYBWCWWWWOB

```

## Step 9

### Thought.

Thought: Great! building the white cross worked and increased the heuristic to 1. I'll continue with the same constraint-based search approach to solve the four white corners and complete the first layer. I'll encode the four down-layer corner constraints, then search short sequences that satisfy each corner's constraints while pruning redundant branches. If at any time the cube becomes solved, I will immediately call final\_answer("") .

2916 L.7 REWARD – AS A GUIDANCE  
2917

Model	Agent Type	Input	Reward	Depth
Claude Sonnet 4	Basic	Full Symbolic	Sticker	2

2921 In this example, the agent explicitly relies on the external reward signal to guide its search. At  
2922 each step, candidate moves are evaluated by their immediate reward: negative rewards trigger an  
2923 immediate undo (e.g.,  $U$  with reward  $-12.0$ ), while positive rewards are retained (e.g.,  $D$  with re-  
2924 ward  $+12.0$ ). The agent’s behavior shows that even a simple reward can provide effective guidance,  
2925 shaping its local search trajectory and preventing it from devolving into purely random exploration.  
2926

2927 *Step 7*

2928 **Thought.**

2929

2930     Thought: Excellent! The L move gave us a +10 reward and we now have a better cube state. Let me check  
2931     the current state and continue looking for more positive moves.  
2932

2933

2934 **Code.**  
2935

2936     # Check current state after the L move  
2937     current\_state = get\_observation()  
2938     print(f"Current state after L: {current\_state}")  
2939  
2940     # Count correct stickers  
2941     correct\_stickers = sum(1 for i in range(54) if current\_state[i] == solved\_state[i])  
2942     print(f"Correct stickers now: {correct\_stickers}/54")  
2943  
2944     # Continue trying more moves to see if we can improve further  
2945     remaining\_moves = ["U", "D", "B"]  
2946     for move in remaining\_moves:  
2947         reward = make\_move(move)  
2948         print(f"Move {move}: reward = {reward}")  
2949  
2950         if reward > 0:  
2951             print(f"Keeping move {move} as it improved the cube")  
2952             break  
2953         else:  
2954             # Undo the move  
2955             opposite\_move = move + "'"  
2956             undo\_reward = make\_move(opposite\_move)  
2957             print(f"Undoing with {opposite\_move}: reward = {undo\_reward}")

2958

2959 **Observation.**  
2960

2961

2962

2963 *Step 8*

2964 **Thought.**

2965

2966

2967     Thought: Great progress! The D move gave us another +12 reward. Let me check the current state and see  
2968     how close we are to solving the cube.  
2969

2970  
2971

## L.8 TRIAL-AND-ERROR TOOL LEARNING

2972  
2973  
2974

Model	Agent Type	Input	Reward	Depth
GPT-5	Standard-Solver	Full Symbolic	No reward	8

2975  
2976  
2977  
2978  
2979  
2980  
2981  
2982  
2983

In this example, the agent initially attempted to convert the environment state into the solver's required format and call the planner (Step 2), but the planner returned errors due to undefined edges. In Steps 3–5, the agent experimented with alternative mappings, but all attempts failed to produce a valid planner input. To overcome this, the agent exhaustively explored all  $4^6$  per-face rotation combinations to find a scheme compatible with the planner's input format, ultimately identifying a scheme compatible with the planner. Upon applying the planner-generated move sequence, the cube was successfully solved. This process exemplifies trial-and-error tool learning, where the agent iteratively experiments, identifies the constraints of the external tool, and adapts its strategy to reliably harness the planner for spatial reasoning tasks.

2984  
2985  
2986

## Step 2

## Thought.

2987  
2988  
2989  
2990  
2991  
2992  
2993  
2994

Thought: I have the current cube state string. Next, I will convert this environment state (color-coded with faces in the order F,B,L,R,U,D) to the solver's required facelet format (URFDLB), using the center colors to map colors to face letters: O->U, Y->R, R->F, W->D, B->L, G->B. I will start with the simplest per-face orientation (top-left to bottom-right) and face order mapping: U(36(44), R(27(35), F(0(8), D(45(53), L(18(26), B(9(17), concatenated in that order. I will call call\_planner with this string and print the solution (or error). If a solution is returned, in the following step I will parse and apply the moves with make\_move, expanding 1/2/3 quarter-turn counts into multiple calls.

2995  
2996

## Code.

2997  
2998  
2999  
3000  
3001  
3002  
3003  
3004  
3005  
3006  
3007  
3008  
3009  
3010  
3011  
3012  
3013  
3014  
3015  
3016  
3017  
3018  
3019  
3020  
3021  
3022  
3023

```

# Build mapping from environment color to solver face letter
color_to_face = {
    'O': 'U', # Up
    'Y': 'R', # Right
    'R': 'F', # Front
    'W': 'D', # Down
    'B': 'L', # Left
    'G': 'B', # Back
}

state = obs # from previous step
# Define the index ranges for each face in the environment (face-on, row-major)
env_faces = {
    'F': list(range(0, 9)),
    'B': list(range(9, 18)),
    'L': list(range(18, 27)),
    'R': list(range(27, 36)),
    'U': list(range(36, 45)),
    'D': list(range(45, 54)),
}

# Solver requires URFDLB order
solver_face_order = ['U', 'R', 'F', 'D', 'L', 'B']

# For now, assume per-face orientation is identity (face-on row-major)
solver_indices = []
for face in solver_face_order:
    if face == 'U':
        solver_indices.extend(env_faces['U'])
    elif face == 'R':
        solver_indices.extend(env_faces['R'])
    elif face == 'F':
        solver_indices.extend(env_faces['F'])
    elif face == 'D':
        solver_indices.extend(env_faces['D'])
    elif face == 'L':
        solver_indices.extend(env_faces['L'])
    elif face == 'B':
        solver_indices.extend(env_faces['B'])

# Convert environment colors to solver face-letter string in the computed order

```

```

3024
3025     try:
3026         solver_str = ''.join(color_to_face[state[i]] for i in solver_indices)
3027     except KeyError as e:
3028         print("Encountered unknown color in state:", e)
3029         solver_str = None
3030
3031     print("Proposed solver string:", solver_str)
3032
3033
3034     if solver_str is not None:
3035         plan = call_planner(tool_cube_representation=solver_str)
3036         print("Planner output:", plan)
3037
3038
3039
3040
3041
3042
3043
3044
3045
3046
3047
3048
3049
3050
3051
3052
3053
3054
3055
3056
3057
3058
3059
3060
3061
3062
3063
3064
3065
3066
3067
3068
3069
3070
3071
3072
3073
3074
3075
3076
3077

```

### Observation.

```

Execution logs:
Proposed solver string: DRFFUDFULBRRDRDBBDFUDFBDFBDRBDULLUFLFLLUURRBURBU
Planner output: Error: Some edges are undefined.
Last output from code snippet:
Error: Some edges are undefined.

```

### Step 6

#### Thought.

```

My previous edge-validation used sorted two-letter strings but compared them to unsorted expectations,
which made all candidates fail.

I will fix the validation by comparing to the alphabetically sorted set of expected edge labels. Then
I'll brute-force face rotations, find a valid scheme, call the planner, and verify the cube is solved.

```

#### Code.

```

# Fix edge validation and brute-force search for a valid per-face rotation scheme, then plan and apply
# moves.

state = obs  # current cube state string

# Environment face indices (row-major as per problem)
env_faces = {
    'F': list(range(0, 9)),
    'B': list(range(9, 18)),
    'L': list(range(18, 27)),
    'R': list(range(27, 36)),
    'U': list(range(36, 45)),
    'D': list(range(45, 54)),
}

def rotate_indices(face_indices, rotation):
    # rotation: 0 (identity), 1 (90° CW), 2 (180°), 3 (270° CW)
    if rotation == 0:
        order = [0,1,2,3,4,5,6,7,8]
    elif rotation == 1:
        order = [6,3,0,7,4,1,8,5,2]
    elif rotation == 2:
        order = [8,7,6,5,4,3,2,1,0]
    elif rotation == 3:
        order = [2,5,8,1,4,7,0,3,6]
    else:
        raise ValueError("Invalid rotation")
    return [face_indices[i] for i in order]

# Build color_to_face mapping from centers
centers = {
    'F': state[4],
    'B': state[9+4],
    'L': state[18+4],
    'R': state[27+4],
}

```

```

3078
3079     'U': state[36+4],
3080     'D': state[45+4],
3081   }
3082   color_to_face = {
3083     centers['U']: 'U',
3084     centers['R']: 'R',
3085     centers['F']: 'F',
3086     centers['D']: 'D',
3087     centers['L']: 'L',
3088     centers['B']: 'B'
3089   }
3090   print("Centers:", centers)
3091   print("Color->Face:", color_to_face)
3092
3093   # Precompute rotated indices for each face and rotation to speed up
3094   pre_rot = {face: {r: rotate_indices(env_faces[face], r) for r in range(4)} for face in
3095   ['U', 'R', 'F', 'D', 'L', 'B']}
3096
3097   def build_solver_str_fast(scheme):
3098     idx = []
3099     idx += pre_rot['U'][scheme['U']]
3100     idx += pre_rot['R'][scheme['R']]
3101     idx += pre_rot['F'][scheme['F']]
3102     idx += pre_rot['D'][scheme['D']]
3103     idx += pre_rot['L'][scheme['L']]
3104     idx += pre_rot['B'][scheme['B']]
3105     return ''.join(color_to_face[state[i]] for i in idx)
3106
3107   def face_base(face_letter):
3108     base = {'U':0, 'R':9, 'F':18, 'D':27, 'L':36, 'B':45}
3109     return base[face_letter]
3110
3111   def get_f(s, face, pos):
3112     return s[face_base(face)+pos]
3113
3114   # Edge definitions in Kociemba URFDLB indexing (face-local 0..8)
3115   edges_def = [
3116     (('U',7), ('F',1)), # UF
3117     (('U',5), ('R',1)), # UR
3118     (('U',1), ('B',1)), # UB
3119     (('U',3), ('L',1)), # UL
3120     (('D',1), ('F',7)), # DF
3121     (('D',5), ('R',7)), # DR
3122     (('D',7), ('B',7)), # DB
3123     (('D',3), ('L',7)), # DL
3124     (('F',5), ('R',3)), # FR
3125     (('F',3), ('L',5)), # FL
3126     (('B',3), ('R',5)), # BR
3127     (('B',5), ('L',3)), # BL
3128   ]
3129
3130   # Expected edges as alphabetically sorted two-letter strings
3131   expected_sorted_edges = set(['FU', 'RU', 'BU', 'LU', 'DF', 'DR', 'BD', 'DL', 'FR', 'FL', 'BR', 'BL'])
3132
3133   def edges_valid(s):
3134     # Build sorted two-letter strings for each edge pair
3135     counts = {}
3136     for (f1,p1), (f2,p2) in edges_def:
3137       a = get_f(s, f1, p1)
3138       b = get_f(s, f2, p2)
3139       pair = ''.join(sorted([a,b]))
3140       if pair not in expected_sorted_edges:
3141         return False
3142       counts[pair] = counts.get(pair, 0) + 1
3143
3144     # Ensure each expected edge appears exactly once
3145     for e in expected_sorted_edges:
3146       if counts.get(e, 0) != 1:
3147         return False
3148     return True
3149
3150   found_scheme = None
3151   found_solver_str = None
3152
3153   # Brute force all 4^6 rotation combinations (4096)
3154   outer_break = False
3155   for u in range(4):
3156     if outer_break: break
3157     for r in range(4):
3158       if outer_break: break
3159       for f in range(4):
3160         if outer_break: break
3161         for d in range(4):
3162           if outer_break: break

```

```

3132
3133     for l in range(4):
3134         if outer_break: break
3135         for b in range(4):
3136             scheme = {'U':u, 'R':r, 'F':f, 'D':d, 'L':l, 'B':b}
3137             s = build_solver_str_fast(scheme)
3138             if edges_valid(s):
3139                 found_scheme = scheme
3140                 found_solver_str = s
3141                 outer_break = True
3142                 break
3143
3144             print("Found scheme:", found_scheme)
3145             print("Solver string preview:", found_solver_str[:30] if found_solver_str else None)
3146
3147             if found_solver_str:
3148                 plan = call_planner(tool_cube_representation=found_solver_str)
3149                 print("Planner output:", plan)
3150
3151
3152
3153
3154
3155
3156
3157
3158
3159
3160
3161
3162
3163
3164
3165
3166
3167
3168
3169
3170
3171
3172
3173
3174
3175
3176
3177
3178
3179
3180
3181
3182
3183
3184
3185

```

### Observation.

```

Centers: {'F': 'R', 'B': 'G', 'L': 'B', 'R': 'Y', 'U': 'O', 'D': 'W'}
Color->Face: {'O': 'U', 'Y': 'R', 'R': 'F', 'W': 'D', 'B': 'L', 'G': 'B'}
Found scheme: {'U': 0, 'R': 0, 'F': 0, 'D': 0, 'L': 0, 'B': 2}
Solver string preview: DRFFUDFULBRDRDBBDFUDFBDBFBBDR
Planner output: R2 L1 B2 R1 U1 R1 B1 U2 D2 F2 L2 F2 U1 D2 (14f)

```

Then executed the solution sequence and solved at Step 7.

## L.9 BYPASSING VISUAL TASKS VIA SYMBOLIC REASONING

We illustrate the performance gap between *Face View* and *Vertex View* inputs using two representative cases.

In the *Face View* setting, the cube is presented as regular  $3 \times 3$  grids. The agent reframed the task as a symbolic parsing problem: mapping each face's colors into the URFDLB string required by the solver. This bypassed visual reasoning, yet worked reliably due to the input's highly ordered structure. The planner accepted the constructed string and returned a valid solution, which the agent executed successfully.

In the *Vertex View* setting, these distorted layouts cannot be parsed directly into solver input without geometric inference. The agent nevertheless attempted symbolic reconstruction via pixel sampling, palette classification, k-means clustering, PCA ordering, and multi-view accumulation, but the resulting cube string was incomplete.

These two cases highlight a key difference between the tasks. In *Face View*, the orderly layout enables a symbolic parsing shortcut that bypasses genuine visual reasoning. In *Vertex View*, this strategy is infeasible, and performance drops sharply. This suggests that even highly capable models tend to prefer symbolic over visual reasoning whenever possible, and that less regular task formats are informative for exposing their limitations in spatial understanding.

### L.9.1 FULL VIEW

Model	Agent Type	Input	Reward	Depth
GPT-5	Ideal-Solver	Full Visual State	no reward	20

Step 7

Thought.

3186  
3187  
3188  
3189  
3190  
3191  
3192  
3193  
3194

Our previous classification mistakenly excluded near-white pixels, undercounting W. I will recompute the sticker colors by averaging patches around each sticker center while only ignoring the gray background and black grid lines (not near-white), then classify to the nearest canonical color. I'll rebuild the 54-character state string, ensure it has 9 of each color, call the planner, expand and apply the moves, and finally submit final\_answer("") .

## Code.

3195  
3196  
3197  
3198  
3199  
3200  
3201  
3202  
3203  
3204  
3205  
3206  
3207  
3208  
3209  
3210  
3211  
3212  
3213  
3214  
3215  
3216  
3217  
3218  
3219  
3220  
3221  
3222  
3223  
3224  
3225  
3226  
3227  
3228  
3229  
3230  
3231  
3232  
3233  
3234  
3235  
3236  
3237  
3238  
3239

```
# Recompute colors including white pixels properly (ignore only gray background and black grid).
img = get_observation()
w, h = img.size
BG = (50,50,50)
BLACK = (0,0,0)

def avg_color_around_include_white(p, r=6):
    cx, cy = p
    samples = []
    for dy in range(-r, r+1):
        y = cy + dy
        if y < 0 or y >= h:
            continue
        for dx in range(-r, r+1):
            x = cx + dx
            if x < 0 or x >= w:
                continue
            rgb = img.getpixel((x,y))
            if rgb == BG or rgb == BLACK:
                continue
            samples.append(rgb)
    if not samples:
        samples = [img.getpixel((cx, cy))]
    sr = sum(p[0] for p in samples)/len(samples)
    sg = sum(p[1] for p in samples)/len(samples)
    sb = sum(p[2] for p in samples)/len(samples)
    return (sr, sg, sb)

canon = {
    'R': (255, 0, 0),
    'G': (0, 255, 0),
    'B': (0, 0, 255),
    'Y': (255, 255, 0),
    'O': (255, 165, 0),
    'W': (255, 255, 255),
}

def classify_avg(rgbf):
    r,g,b = rgbf
    best=None; bd=1e18
    for k,v in canon.items():
        d=(r-v[0])**2+(g-v[1])**2+(b-v[2])**2
        if d<bd:
            bd=d; best=k
    return best

def colors_from_grid_patch_inc_white(grid_points):
    return [classify_avg(avg_color_around_include_white(p, r=6)) for p in grid_points]

# Use previously computed grids U_grid, F_grid, D_grid, B_grid, L_grid, R_grid
U_colors = colors_from_grid_patch_inc_white(U_grid)
F_colors = colors_from_grid_patch_inc_white(F_grid)
D_colors = colors_from_grid_patch_inc_white(D_grid)
B_colors = colors_from_grid_patch_inc_white(B_grid)
L_colors = colors_from_grid_patch_inc_white(L_grid)
R_colors = colors_from_grid_patch_inc_white(R_grid)

from collections import Counter
state = ''.join(F_colors + B_colors + L_colors + R_colors + U_colors + D_colors)
print("Color counts:", Counter(state), "length:", len(state))
print("State:", state)

# Call planner
plan = call_golden_planner(tool_cube_representation=state)
print("Planner:", plan)

# Parse and expand moves like 'L3 U1 B1 R2 ... (19f)'
import re
tokens = re.findall(r'([URFDLB]) ([123])', plan)
```

```

3240
3241     expanded = []
3242     for face, num in tokens:
3243         n = int(num)
3244         if n == 1:
3245             expanded.append(face)
3246         elif n == 2:
3247             expanded.extend([face, face])
3248         elif n == 3:
3249             expanded.extend([face, face, face])
3250     print("Expanded moves:", expanded)
3251
3252
3253
3254     # Apply moves
3255     for mv in expanded:
3256         print(make_move(move=mv))
3257
3258     # Submit final check
3259     final_answer("")

```

### Observation.

```

3256
3257     Color counts: Counter({'W': 9, 'Y': 9, 'R': 9, 'B': 9, 'O': 9, 'G': 9}) length: 54
3258     State: WYYRBOWGYWRWGRWGOOYGBRWYOYOOORYOBGBBYBROGGGRGBWYBRRW
3259     Planner: U3 F2 R2 U2 F2 D2 F2 L1 B3 R1 U2 F3 D3 L3 B3 U2 L2 R1 D3 U1 (20f)
3260     Expanded moves: ['U', 'U', 'U', 'F', 'F', 'R', 'R', 'U', 'U', 'F', 'F', 'D', 'D', 'F', 'F', 'L', 'B',
3261     'B', 'B', 'R', 'U', 'U', 'F', 'F', 'F', 'D', 'D', 'D', 'L', 'L', 'B', 'B', 'U', 'U', 'L',
3262     'L', 'R', 'D', 'D', 'D', 'U']
3263     Applied move 'U' on the cube
3264     Applied move 'U' on the cube
3265     Applied move 'U' on the cube
3266     Applied move 'F' on the cube
3267     Applied move 'F' on the cube
3268     Applied move 'R' on the cube
3269
3270
3271     // ...
3272     Applied move 'R' on the cube
3273     Applied move 'D' on the cube
3274     Applied move 'D' on the cube
3275     Applied move 'D' on the cube
3276     Applied move 'U' on the cube

```

Solved.

### L.9.2 VERTEX VIEW

Model	Agent Type	Input	Reward	Depth
GPT-5	Standard-Solver	Vertex View	no reward	20

Step 7

Thought.

```

3282
3283     The previous approach hit the operation limit due to heavy k-means and broad view sweeps, and also
3284     misclassified colors. I'll switch to a lighter pipeline:
3285     - Classify colors by nearest reference palette (O/W/R/G/B/Y) for robustness.
3286     - Sample far fewer pixels (stride=5) and run k-means with fewer iterations (k=27, iter=8) to find
3287     sticker centers in each view.
3288     - Group the 27 clusters spatially into 3 faces, order each face's 3x3 via PCA, and use the center
3289     sticker to name the face.
3290     - Collect exactly 5 views (right, right, down, right, right) to cover all faces.
3291     - Convert the six color grids into the URFDLB string (O,Y,R,W,B,G mapped to U,R,F,D,L,B), call the
3292     planner, apply the moves, then check solved state.
3293
3294     If the planner errors (mapping inconsistent), I'll print the composed string to debug. Otherwise I'll
3295     execute the plan and finish.

```

Code.

```

3294
3295
3296
3297
3298
3299
3300
3301
3302
3303
3304
3305
3306
3307
3308
3309
3310
3311
3312
3313
3314
3315
3316
3317
3318
3319
3320
3321
3322
3323
3324
3325
3326
3327
3328
3329
3330
3331
3332
3333
3334
3335
3336
3337
3338
3339
3340
3341
3342
3343
3344
3345
3346
3347
3348
3349
3350
3351
3352
3353
3354
3355
3356
3357
3358
3359
3360
3361
3362
3363
3364
3365
3366
3367
3368
3369
3370
3371
3372
3373
3374
3375
3376
3377
3378
3379
3380
3381
3382
3383
3384
3385
3386
3387
3388
3389
3390
3391
3392
3393
3394
3395
3396
3397
3398
3399
3400
3401
3402
3403
3404
3405
3406
3407
3408
3409
3410
3411
3412
3413
3414
3415
3416
3417
3418
3419
3420
3421
3422
3423
3424
3425
3426
3427
3428
3429
3430
3431
3432
3433
3434
3435
3436
3437
3438
3439
3440
3441
3442
3443
3444
3445
3446
3447
3448
3449
3450
3451
3452
3453
3454
3455
3456
3457
3458
3459
3460
3461
3462
3463
3464
3465
3466
3467
3468
3469
3470
3471
3472
3473
3474
3475
3476
3477
3478
3479
3480
3481
3482
3483
3484
3485
3486
3487
3488
3489
3490
3491
3492
3493
3494
3495
3496
3497
3498
3499
3500
3501
3502
3503
3504
3505
3506
3507
3508
3509
3510
3511
3512
3513
3514
3515
3516
3517
3518
3519
3520
3521
3522
3523
3524
3525
3526
3527
3528
3529
3530
3531
3532
3533
3534
3535
3536
3537
3538
3539
3540
3541
3542
3543
3544
3545
3546
3547
3548
3549
3550
3551
3552
3553
3554
3555
3556
3557
3558
3559
3560
3561
3562
3563
3564
3565
3566
3567
3568
3569
3570
3571
3572
3573
3574
3575
3576
3577
3578
3579
3580
3581
3582
3583
3584
3585
3586
3587
3588
3589
3590
3591
3592
3593
3594
3595
3596
3597
3598
3599
3600
3601
3602
3603
3604
3605
3606
3607
3608
3609
3610
3611
3612
3613
3614
3615
3616
3617
3618
3619
3620
3621
3622
3623
3624
3625
3626
3627
3628
3629
3630
3631
3632
3633
3634
3635
3636
3637
3638
3639
3640
3641
3642
3643
3644
3645
3646
3647
3648
3649
3650
3651
3652
3653
3654
3655
3656
3657
3658
3659
3660
3661
3662
3663
3664
3665
3666
3667
3668
3669
3670
3671
3672
3673
3674
3675
3676
3677
3678
3679
3680
3681
3682
3683
3684
3685
3686
3687
3688
3689
3690
3691
3692
3693
3694
3695
3696
3697
3698
3699
3700
3701
3702
3703
3704
3705
3706
3707
3708
3709
3710
3711
3712
3713
3714
3715
3716
3717
3718
3719
3720
3721
3722
3723
3724
3725
3726
3727
3728
3729
3730
3731
3732
3733
3734
3735
3736
3737
3738
3739
3740
3741
3742
3743
3744
3745
3746
3747
3748
3749
3750
3751
3752
3753
3754
3755
3756
3757
3758
3759
3760
3761
3762
3763
3764
3765
3766
3767
3768
3769
3770
3771
3772
3773
3774
3775
3776
3777
3778
3779
3780
3781
3782
3783
3784
3785
3786
3787
3788
3789
3790
3791
3792
3793
3794
3795
3796
3797
3798
3799
3800
3801
3802
3803
3804
3805
3806
3807
3808
3809
3810
3811
3812
3813
3814
3815
3816
3817
3818
3819
3820
3821
3822
3823
3824
3825
3826
3827
3828
3829
3830
3831
3832
3833
3834
3835
3836
3837
3838
3839
3840
3841
3842
3843
3844
3845
3846
3847
3848
3849
3850
3851
3852
3853
3854
3855
3856
3857
3858
3859
3860
3861
3862
3863
3864
3865
3866
3867
3868
3869
3870
3871
3872
3873
3874
3875
3876
3877
3878
3879
3880
3881
3882
3883
3884
3885
3886
3887
3888
3889
3890
3891
3892
3893
3894
3895
3896
3897
3898
3899
3900
3901
3902
3903
3904
3905
3906
3907
3908
3909
3910
3911
3912
3913
3914
3915
3916
3917
3918
3919
3920
3921
3922
3923
3924
3925
3926
3927
3928
3929
3930
3931
3932
3933
3934
3935
3936
3937
3938
3939
3940
3941
3942
3943
3944
3945
3946
3947
3948
3949
3950
3951
3952
3953
3954
3955
3956
3957
3958
3959
3960
3961
3962
3963
3964
3965
3966
3967
3968
3969
3970
3971
3972
3973
3974
3975
3976
3977
3978
3979
3980
3981
3982
3983
3984
3985
3986
3987
3988
3989
3990
3991
3992
3993
3994
3995
3996
3997
3998
3999
3999

```

```

3348
3349     if xs:
3350         newc.append((sum(xs)/len(xs), sum(ys)/len(ys)))
3351     else:
3352         newc.append(centers[j])
3353     centers = newc
3354     clusters = [[] for _ in range(k)]
3355     for i in range(len(points)):
3356         clusters[assigns[i]].append(i)
3357     return centers, assigns, clusters
3358
3359     def majority_color(color_list):
3360         counts = defaultdict(int)
3361         for rgb in color_list:
3362             c = rgb_to_color_letter(rgb)
3363             counts[c] += 1
3364         return max(counts.items(), key=lambda kv: kv[1])[0] if counts else 'W'
3365
3366     def extract_27_stickers(img):
3367         pts, cols = sample_colored_pixels(img, stride=5, max_samples=1500)
3368         print("Sampled colored:", len(pts))
3369         res = kmeans_points(pts, 27, iters=8)
3370         if res[0] is None:
3371             return []
3372         _, assigns27, clusters27 = res
3373         stickers = []
3374         for j in range(27):
3375             idxs = clusters27[j]
3376             if not idxs:
3377                 continue
3378             cx = sum(pts[i][0] for i in idxs)/len(idxs)
3379             cy = sum(pts[i][1] for i in idxs)/len(idxs)
3380             col = majority_color([cols[i] for i in idxs])
3381             stickers.append((cx, cy, col))
3382
3383         return stickers
3384
3385     def pca_axes(points):
3386         if not points:
3387             return (1.0,0.0), (0.0,1.0)
3388         mx = sum(p[0] for p in points)/len(points)
3389         my = sum(p[1] for p in points)/len(points)
3390         sxx = sum((p[0]-mx)**2 for p in points)/len(points)
3391         syy = sum((p[1]-my)**2 for p in points)/len(points)
3392         sxy = sum((p[0]-mx)*(p[1]-my) for p in points)/len(points)
3393         tr = sxx+syy
3394         det = sxx*syy - sxy*sxy
3395         disc = max(tr*tr/4 - det, 0.0)
3396         lam1 = tr/2 + sqrt(disc)
3397         if abs(sxy) + abs(lam1 - sxx) < 1e-9:
3398             u=(1.0,0.0)
3399         else:
3400             u=(sxy, lam1 - sxx)
3401             n = sqrt(u[0]*u[0]+u[1]*u[1])
3402             u=(u[0]/n, u[1]/n)
3403             v=(-u[1], u[0])
3404         return u,v
3405
3406     def kmeans_scalar(values, k=3, iters=8):
3407         # 1D kmeans implemented via 2D wrapper to reuse
3408         pts = [(v,0.0) for v in values]
3409         _, assigns, clusters = kmeans_points(pts, k, iters)
3410         return assigns, clusters
3411
3412     def group_into_faces(stickers):
3413         centers=[s[0] for s in stickers]
3414         res = kmeans_points(centers, 3, iters=8)
3415         if res[0] is None:
3416             return []
3417         _, assigns3, clusters3 = res
3418         faces=[]
3419         for j in range(3):
3420             idxs = clusters3[j]
3421             group = [stickers[i] for i in idxs]
3422             if len(group) >= 7:
3423                 faces.append(group)
3424
3425         return faces
3426
3427     def order_face_grid(face_stickers):
3428         pts=[s[0] for s in face_stickers]
3429         u,v = pca_axes(pts)
3430         proj=[]
3431         for i,(pt,col) in enumerate(face_stickers):
3432             pu = pt[0]*u[0] + pt[1]*u[1]
3433             pv = pt[0]*v[0] + pt[1]*v[1]

```

```

3402
3403     proj.append((pu,pv,i))
3404     col_vals = [p[0] for p in proj]
3405     row_vals = [p[1] for p in proj]
3406     cols_assigns, _ = kmeans_scalar(col_vals, 3, iters=8)
3407     rows_assigns, _ = kmeans_scalar(row_vals, 3, iters=8)
3408     col_groups = defaultdict(list); row_groups = defaultdict(list)
3409     for idx,(pu,pv,i) in enumerate(proj):
3410         col_groups[cols_assigns[idx]].append(i)
3411         row_groups[rows_assigns[idx]].append(i)
3412     ordered_cols = sorted(col_groups.keys(), key=lambda c: sum(face_stickers[i][0][0] for i in col_groups[c])/len(col_groups[c]))
3413     ordered_rows = sorted(row_groups.keys(), key=lambda r: sum(face_stickers[i][0][1] for i in row_groups[r])/len(row_groups[r]))
3414     grid = [[None]*3 for _ in range(3)]
3415     for idx,(pu,pv,i) in enumerate(proj):
3416         rlab = rows_assigns[idx]; clab = cols_assigns[idx]
3417         r = ordered_rows.index(rlab); c = ordered_cols.index(clab)
3418         grid[r][c] = face_stickers[i][1]
3419     return grid
3420
3421
3422
3423
3424
3425
3426
3427
3428
3429
3430
3431
3432
3433
3434
3435
3436
3437
3438
3439
3440
3441
3442
3443
3444
3445
3446
3447
3448
3449
3450
3451
3452
3453
3454
3455
def faces_from_view(img):
    stickers = extract_27_stickers(img)
    print("Sticker clusters:", len(stickers))
    faces_groups = group_into_faces(stickers)
    print("Faces in this view:", len(faces_groups))
    faces = {}
    for group in faces_groups:
        grid = order_face_grid(group)
        center = grid[1][1]
        faces[center] = {"grid": grid}
    return faces

def accumulate_faces(acc, new_faces):
    for k,v in new_faces.items():
        if k not in acc:
            acc[k]=v

# Collect faces from a fixed small sequence of views
observed = {}
img = get_observation()
accumulate_faces(observed, faces_from_view(img))
apply_view_transformation("view_right")
img = get_observation()
accumulate_faces(observed, faces_from_view(img))

print("Collected faces:", list(observed.keys()))
for k in observed:
    print(k, observed[k]["grid"])

# Map to solver's URFDLB letters
color_to_face_letter = {'O':'U', 'W':'D', 'R':'F', 'G':'B', 'B':'L', 'Y':'R'}

def grid_to_str(grid):
    return ''.join(''.join(grid[r][c] for c in range(3)) for r in range(3))

def compose_URFDLB(obs):
    # URFDLB color centers order: O (U), Y (R), R (F), W (D), B (L), G (B)
    order = ['O','Y','R','W','B','G']
    s=""
    for ck in order:
        if ck not in obs:
            print("Warning missing face:", ck)
            grid = [[ck]*3 for _ in range(3)]
        else:
            grid = obs[ck]["grid"]
            mapped = [[color_to_face_letter.get(ch, 'U') for ch in row] for row in grid]
            s += grid_to_str(mapped)
    return s

solver_str = compose_URFDLB(observed)
print("Solver string:", solver_str)

plan = call_planner(tool_cube_representation=solver_str)
print("Planner output:", plan)

```

```

3456
3457
3458     def execute_plan(plan_str):
3459         toks = plan_str.strip().split()
3460         for t in toks:
3461             if '(' in t: # end marker with "(xf"
3462                 break
3463             face = t[0]
3464             num = t[1] if len(t) > 1 else '1'
3465             if num == '1':
3466                 print(make_move(face))
3467             elif num == '2':
3468                 print(make_move(face))
3469             print(make_move(face))
3470             elif num == '3':
3471                 print(make_move(face + ""))
3472
3473     if plan and not plan.strip().startswith("Error"):
3474         execute_plan(plan)
3475
3476     # Final check
3477     final_answer("")

```

### Observation.

```

3473     Sampled colored: 106
3474     Sticker clusters: 27
3475     Faces in this view: 3
3476     Sampled colored: 106
3477     Sticker clusters: 27
3478     Faces in this view: 3
3479     Sampled colored: 106
3480     Sticker clusters: 27
3481     Faces in this view: 3
3482     Sampled colored: 103
3483     Sticker clusters: 27
3484     Faces in this view: 2
3485     Sampled colored: 102
3486     Sticker clusters: 27
3487     Faces in this view: 2
3488     Collected faces: ['B', 'R', None, 'O', 'G', 'W']
3489     B [['Y', 'Y', 'O'], ['O', 'B', 'O'], [None, 'B', None]]
3490     R [[None, 'Y', 'W'], ['R', 'R', 'Y'], [None, None, 'B']]
3491     None [[R', None, 'G'], [W', None, None], [G', 'Y', 'Y']]
3492     O [[R', 'R', 'O'], [G', 'O', 'B'], [None, 'G', None]]
3493     G [[B', 'G', 'W'], ['R', 'G', 'W'], [R', 'O', None]]
3494     W [[W', 'W', 'Y'], ['Y', 'W', 'W'], [B', 'Y', None]]
3495     Warning missing face: Y
3496     Solver string: FFUBULUBURRRRRRRURDFRUULDDRRDRLURRUULUULULBDFBDFUU
3497     Planner output: Error: Cube definition string FFUBULUBURRRRRRRURDFRUULDDRRDRLURRUULUULULBDFBDFUU
3498     does not contain exactly 9 facelets of each color.

```

### Not Solved.

```

3499
3500
3501
3502
3503
3504
3505
3506
3507
3508
3509

```