

## 478 **Appendix**

479 This appendix provides detailed information about our method and experimental setup. It is organized  
480 as follows:

- 481 • In section A, we first describe the prompting scheme used by different agents in InstructFlow,  
482 including both the shared initialization prompt and agent-specific templates for the planner,  
483 code generator, and symbolic constraint generator.
- 484 • In section B.1, we provide additional details on our experimental setup, including how  
485 programs are instantiated and executed within the simulation pipeline. We also identify  
486 critical limitations in the original PRoC3S evaluation protocol and introduce a VLM-based  
487 semantic check to address them.
- 488 • In section B.2.1 , we present an extended experiment exploring the effect of visual inputs  
489 on planning, followed by a summary of symbolic constraints induced across different  
490 manipulation tasks in section B.2.2.
- 491 • In section B.2.3, we include several representative case studies that illustrate how In-  
492 structFlow recovers from planning failures through symbolic reasoning and instruction  
493 graph-guided code repair.

## 494 A InstructFlow Prompting Details

495 Here we provide details on the prompting scheme used for each agent in InstructFlow. The prompting  
496 template consists of two parts: a shared initialization template and agent-specific prompts. The  
497 complete structure is illustrated as follows:

498 Shared Prompt Templates:

```
499 1 {{{system_prompt}}}
500 2 {{{domain_setup_code}}}
501 3 {{{skill_preface}}}
502 4 {{{domain_skills}}}
```

503 Agent-Specific Prompt Templates:

```
504 1 {{{planner_role}}}
505 2 {{{code_generator_role}}}
506 3 {{{constraint_generator_role}}}
```

### 507 A.1 Shared Prompt Templates

508 All agents share a common initial prompt structure, consisting of: (a) `system_prompt`, (b)  
509 `domain_setup_code`, (c) `skill_preface`, (d) `domain_skills`. This shared context is constructed  
510 following the initial prompt setup introduced in PRoC3S [3], where detailed environment and skill  
511 specifications are defined. We adopt the same structure without modification, and refer readers to the  
512 original paper for full specification details.

### 513 A.2 Agent-Specific Prompt Templates

514 As outlined in Section 3, we have three primary agents: the **InstructFlow Planner**, which decom-  
515 poses the high-level goal into structured subgoals through instruction graph construction; the **Code**  
516 **Generator**, which translates the instruct graph into executable code snippets; and the **Constraint**  
517 **Generator**, which analyzes execution feedback to induce symbolic constraints for graph-guided code  
518 repair.

#### 519 A.2.1 InstructFlow Planner

##### Instruction Graph Construction Prompt

```
<inputs>
  {task_goal}
  {initial_state}
</inputs>
```

You are responsible for constructing an adaptive instruction graph that serves as an intermediate reasoning structure for robotic task planning. Your task is to generate a sequence of base planning nodes that decompose the task into subgoals. These nodes form the initial executable backbone of the plan. Do not include any reasoning nodes at this stage. Each node should include a semantic description of the subgoal, which can represent either:

- a concrete manipulation action (e.g., pick, place), or
- a prerequisite operation (e.g., selecting a target object, removing an obstacle).

Ensure subgoals are:

- logically ordered,
- collectively sufficient to achieve the goal.

*# Node Type: Planning Node*

```
{ "id": "n1",
  "type": "action",
  "inputs": ["goal", "initial_state"],
  "output": <"natural language subgoal">}
```

520

## Instruction Graph Revision Prompt

```
<inputs>
  {task_goal}
  {initial_state}
  {last_instruction_graph}
  {symbolic_constraint}
</inputs>
```

Your task is to revise the instruction graph in response to symbolic constraints inferred from execution failures. This process involves constructing a reasoning flow that incrementally generates a refined instruction to update downstream planning nodes. You need to follow these steps:

### Step 1: Reasoning Node Selection

You need to insert one or more reasoning nodes upstream of affected planning nodes guiding by symbolic constraints. Each reasoning node should perform a symbolic transformation relevant to the failure, and output an intermediate instruction fragment.

### Step 2: Instruction Flow Construction

Sequentially process each reasoning node, using its output as an instruction fragment that incrementally updates the evolving task description. These fragments form a directed information flow. At the end of the reasoning node, concatenate all instruction fragments to form a single composite instruction that encodes the full reasoning chain and can be used to update downstream planning nodes.

### Step 3: Instruction Graph Update

Use the composite instruction to update the planning nodes. Replace the original subgoal with this revised version, which incorporates both the task intent and refinements. The updated subgoals of instruction graph should then be passed as the final instruction to the code generator.

We define five types of reasoning nodes, each with its specific functionality and structure as described below:

-The spatial relation reasoning node analyzes pairwise spatial relationships between all visible objects to generate a spatial relation graph:

*# Node Type: Spatial Relation Reasoning Node*

```
{ "id": "n2",
  "type": "spatial_relation_reasoning_node",
  "input": ["initial_state"],
  "output": <"spatial_relation_graph">}
```

-The object density reasoning node estimate the local spatial density around each object to reflect how crowded its surroundings are.

*# Node Type: Object Density Reasoning Node*

```
{ "id": "n3",
  "type": "object_density_analysis_node",
  "input": ["initial_state"],
  "output": <"object_density_map">}
```

-The object selection reasoning node combines the goal, the output of the spatial perception node, the output of the object density analysis node as the inputs to infer which correct manipulated object(s) should be selected to accomplish the goal.

*# Node Type: Object Selection Reasoning Node*

```
{ "id": "n4",
  "type": "object_selection_reasoning_node",
  "input": ["goal", "spatial_relations_graph", "object_density_map", "symbolic_constraint"],
  "output": <"manipulated_objects">}
```

-The plan logic Reasoning Node combines the goal, the output of object selection reasoning node, and the symbolic predicate as the inputs to infer the correct execution order among the manipulated targets.

*# Node Type: Plan Logic Reasoning Node*

```
{ "id": "n5",
  "type": "plan_logic_reasoning_node",
  "input": ["goal", "manipulated_objects", "symbolic_constraint"],
  "output": <"execution_order">}
```

-The parameter range reasoning node combines the goal and the symbolic predicate to adjusts the ranges of action plan parameters based on explicit rules defined in symbolic predicates, and outputs instructions indicating whether to expand or shrink the ranges to meet task requirements.

*# Node Type: Parameter Range Reasoning Node*

```
{ "id": "n6",
  "type": "parameter_range_adjustment_node",
  "inputs": ["goal", "symbolic_constraint"],
  "output": <"range_adjustment_instruction">}
```

## 522 A.2.2 Code Generator

### Code Generator Prompt

```
<inputs>
  {task_goal}
  {initial_state}
  {subgoals_instruction}
</inputs>
```

You are a code generation agent in a robotic planning system. Your goal is to generate two things:

First, generate a python function named 'gen\_plan' that can take any discrete or continuous inputs. No list inputs are allowed and return the entire plan with all steps included where the parameters to the plan depend on the inputs. The plan should be generated based on the initial high-level computation graph, which is composed of a sequence of subgoals. Each subgoal corresponds to either a manipulation action or a prerequisite operation.

Second, generate a python function 'gen\_domain' that returns a set of bounds for the continuous or discrete input parameters. The number of bounds in the generated domain should exactly match the number of inputs to the function excluding the state input.

The function you give should always achieve the goal regardless of what parameters from the domain are passed as input. The 'gen\_plan' function therefore defines a family of solutions to the problem. Explain why the function will always satisfy the goal regardless of the input parameters. Make sure your function inputs allow for as much variability in output plan as possible while still achieving the goal. Your function should be as general as possible such that any correct answer corresponds to some input parameters to the function.

The main function should be named EXACTLY 'gen\_plan' and the domain of the main function should be named EXACTLY 'gen\_domain'. Do not change the names. Do not create any additional classes or overwrite any existing ones. Aside from the initial state all inputs to the 'gen\_plan' function MUST NOT be of type List or Dict. List and Dict inputs to 'gen\_plan' are not allowed. Additionally, the input to 'gen\_domain' must be exactly the 'initial:RavenBelief' argument, even if this isn't explicitly used within the function!

## 525 A.2.3 Symbolic Constraint Generator

### Constraint Generator Prompt

```
<inputs>
  {task_goal}
  {initial_state}
  {failure_feedback}
  {generated_code}
</inputs>
```

You are a symbolic reasoning agent tasked with diagnosing execution failures in robotic manipulation tasks. Your goal is to induce generalizable symbolic constraints that explain the failure and can guide future plan correction. Your task is to perform the following reasoning steps to generate symbolic constraint(s) for plan repair:

Step 1. Failure-relevant retrieval:

Identify the exact code segment responsible for triggering the failure. In a failure feedback, "Step N" refers to the N-th action generated by the code (zero-based index). Specifically, it corresponds to the N-th `plan.append(Action(...))` call in the code. For example: "Step 1, Action: place" refers to the second action in the plan (index = 1).

Step 2. Code-level reasoning:

Extract all variables relevant to this failure, including:

- Manipulated objects.
- Object poses and spatial relations.
- Action parameters (e.g., offsets).
- Domain ranges for parameters.

Step 3. Diagnostic reasoning:

Based on the environment state and extracted variables, analyze the geometric or physical cause of failure. You must consider multiple possible causes, including but not limited to:

- Geometric violations (e.g., collisions, unstable placement, path obstruction).
- Temporal inconsistencies (e.g., incorrect subgoal ordering, premature actions).
- Symbolic logical errors (e.g., wrong object selection, missing prerequisite conditions).

For physical causes, compute diagnostic metrics, including but not limited to:

- Proximity distances between relevant entities.
- Existence of collision-free paths.
- Stability metrics (e.g., center of mass projection).

For symbolic or logical causes, analyze:

- Whether the current action respects task-specific symbolic constraints.

- Dependencies between subgoals as defined in the instruction graph.
- Whether preconditions for the current subgoal are satisfied.

Provide a concise diagnosis explaining why the failure occurred, explicitly stating:

- The type of cause (physical, temporal, symbolic).
- The reasoning process leading to this conclusion.

You are not allowed to assume the cause based solely on the failure description. All conclusions must be verified through concrete reasoning over environment state, task semantics, or diagnostic metrics.

Step 4. Symbolic Constraint Induction:

Given the diagnosis and variables, formulate a symbolic constraint that abstracts the failure into a general, reusable rule. You must express the constraint using standard symbolic predicates, including but not limited to:

- Spatial relations: On, ClearOf, Aligned, StableOn...
- Temporal/ordering: Before, Precondition, Order...
- Semantic affordances: Reachable, Occludes, Graspable, Affords...
- Physical feasibility: ProximitySafe, PathClear, PlacementFeasible, ForceStable...

The constraint must express logical conditions involving task entities, their relationships, or parameter ranges, and be parameterized with thresholds or bounds when applicable.

527

## 528 **B Experiment Details and Additional Results**

### 529 **B.1 Experimental Setup Details**

#### 530 **B.1.1 Program Instantiation and Simulation Pipeline**

531 Our framework follows the same two-stage code generation and execution process used in PRoC3S [3].  
532 Specifically, after generating the code that defines a task plan via LLM, the resulting Python function  
533 consists of two components: `get_plan()` and `get_domain()`.

534 The `get_plan()` function encodes a sequence of symbolic actions with continuous or discrete  
535 parameters, corresponding to the subgoals decomposed in the instruction graph. The `get_domain()`  
536 function specifies the sampling bounds for each parameter from a predefined sampler. These samplers  
537 (e.g., `ContinuousSampler`, `GraspSampler`) generate candidate values for plan parameters without  
538 awareness of environmental constraints such as collisions or instability.

539 To ground the abstract plan into an executable one, we adopt the same strategy as PRoC3S: sample  
540  $n$  parameter instantiations from the domains defined in `get_domain()`, and for each instantiation,  
541 evaluate the resulting plan in a physics-based simulator. If the plan violates any constraints (e.g.,  
542 Kinematic, collisions, grasp, placement constraints), the simulator reports detailed constraint violation  
543 feedback. Once a constraint-free plan is found in simulation, it is deployed in the real environment.

#### 544 **B.1.2 Fixing Protocol Limitations in PRoC3S Evaluation**

545 While our experimental setup builds directly on the original PRoC3S [3] framework and reuses  
546 its environment, skill library, and simulation interface, we identified structural limitations in its  
547 evaluation protocol. Specifically, PRoC3S treats any execution that does not explicitly violate  
548 simulator constraints as successful, regardless of whether the task goal has been semantically achieved.

549 **Incorrect but constraint-free Plan Misclassified as Successful** For example, a plan may place an  
550 object in an incorrect position, fail to form the required structure (e.g., a pyramid), or manipulate  
551 the wrong object altogether. As long as no collisions or instability are triggered in simulation, such  
552 plans are incorrectly classified as successful. Figure 6 illustrate several cases where the task goal was  
553 clearly unmet, yet no feedback was generated to initiate replanning.

554 **Fixing the Evaluation Protocol via VLM Check** To address this fundamental evaluation gap, we  
555 introduce a semantic-level verification step using a vision-language model (GPT-4o). After executing  
556 each plan, we render the final scene and prompt the VLM to assess whether the natural language  
557 goal has been achieved. If not, the failure is recorded and propagated, triggering a replanning  
558 cycle. **Specifically, we explicitly instruct the VLM to evaluate structured visual conditions and  
559 provide clear success criteria along with positive and negative examples, enabling it to produce  
560 consistent and grounded success judgments.**

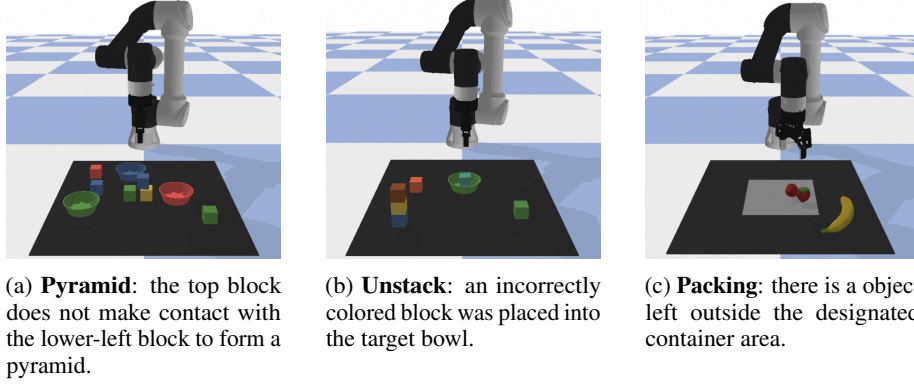


Figure 6: Task failures caused by incorrect plans that are accepted as successful by PRoC3S, despite violating the task goal, because they do not trigger any constraint violations.

561 This fix does not modify the core PRoC3S planning mechanism, but augments the evaluation  
 562 logic with a reliable, goal-aware success signal. It enables all baselines—including PRoC3S and  
 563 InstructFlow—to be assessed under a consistent, semantically meaningful criterion. While this change  
 564 may lower the reported success rates of prior methods, we consider it essential for fair and rigorous  
 565 comparison, especially in tasks with under-specified goals or ambiguous execution semantics.

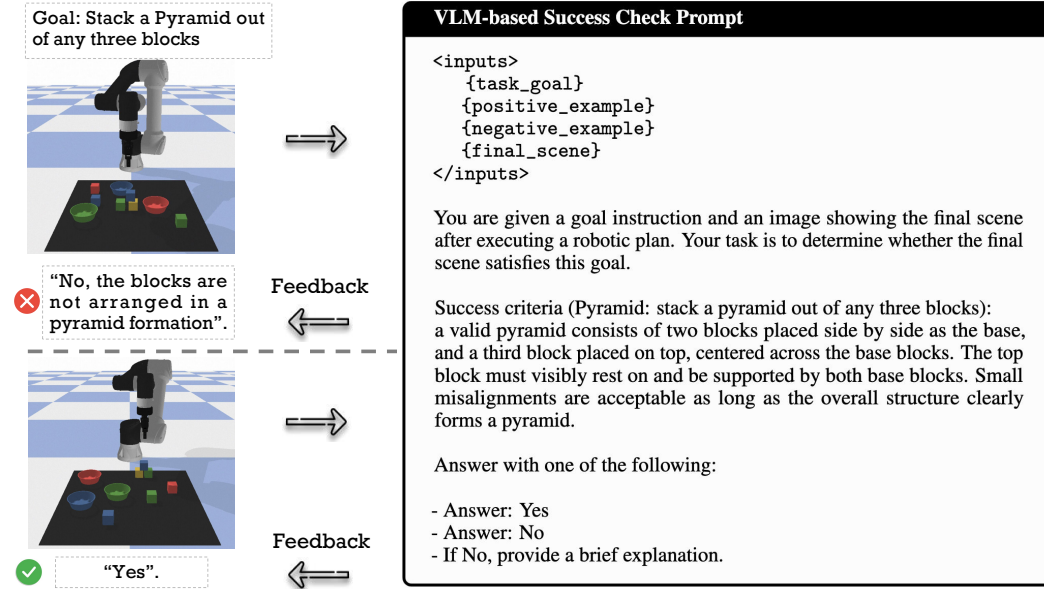


Figure 7: VLM-based success check on two execution plans for the Pyramid task. The upper plan produces a **non-standard pyramid** where the top block does not contact the left base block. The VLM detects this semantic error and returns “No” with reason. The lower plan yields a **correct pyramid** structure, and the VLM responds with “Yes”. This demonstrates the importance of incorporating VLM check as a complementary layer to evaluation protocol.

## 566 B.2 Additional Experiment Results and Cases

### 567 B.2.1 Additional Experiment

568 To assess the potential of vision-language models (VLMs) for improving task planning, we conduct  
 569 an additional experiment where our existing GPT-4o-based reasoning agent receive both the goal and  
 570 an image of the initial scene, enabling multimodal visual reasoning.

This vision-augmented setup improves the planner’s ability to reason under partial observability or ambiguous symbolic states. As shown in Table 3, we observe performance gains in several visually grounded tasks, including *Packing (Blocks)*, *Unstack*, and *Packing (YCB)*, where early-stage visual cues help the system make better object selection or ordering decisions.

	Drawing				Arrange Blocks				Arrange YCB	
	Star	Arrow	Letters	Enclosed	Pyramid	Line	Packing	Unstack	Packing	Stacking
LLM <sup>3</sup> (w/ vision)	40%	40%	80%	50%	0%	40%	30%	0%	0%	10%
CaP (w/ vision)	10%	0%	40%	30%	20%	20%	20%	20%	40%	10%
PRoC3S (w/ vision)	90%	<b>80%</b>	80%	90%	60%	70%	60%	80%	40%	40%
InstructFlow (w/ vision)	<b>100%</b>	<b>80%</b>	<b>100%</b>	<b>100%</b>	<b>90%</b>	<b>100%</b>	<b>90%</b>	<b>100%</b>	<b>60%</b>	<b>70%</b>

Table 3: Task success rates (%) across drawing, block arrangement, and YCB manipulation tasks using VLM-based visual reasoning. Bold indicates the best-performing results.

## B.2.2 Table of Symbolic Constraints Discovered

Table 4 summarizes the symbolic constraints that were automatically induced based on failure feedback across different manipulation tasks. From a content perspective, we highlight the following key properties:

**Structural Consistency.** All induced constraints conform to the symbolic constraint formulation presented in Section 3.2. Each constraint instance can be expressed either as a relational predicate  $R_i(e_a, e_b)$  or a physical threshold condition  $f_j(\Theta) \oplus \tau$ . This ensures that all constraints are logically composable, interpretable, and grounded in the formal symbolic space defined by  $\phi$ .

**Diverse Coverage of Constraint Types.** The constraint set spans a wide range of task-relevant constraint categories, including:

- **Spatial safety:** ProximitySafe, ClearOf, PathClear...
- **Placement feasibility:** PlacementFeasible, Aligned, StableOn...
- **Geometric parameters:** Distance, Offset, CenterOfMass...
- **Temporal logic:** Before, Order

These categories capture both physical feasibility and symbolic reasoning failure modes, supporting a broad range of corrective strategies.

Notably, the table does not include constraints for the Drawing tasks. This is because the drawing tasks are comparatively simpler in structure and were typically solved by the planner in a single attempt without triggering any failure-driven refinement. As a result, no symbolic constraint induction process was invoked for these tasks, and they are therefore excluded from this table.

## B.2.3 Case study

In addition to the *Unstack* task discussed in section 4.3, we present several representative manipulation tasks as case studies to further demonstrate the effectiveness and generality of our method. For each task, we analyze the core planning difficulties, the symbolic constraints induced during execution, and how InstructFlow dynamically updates the instruction graph to recover from failures.

**Pyramid.** The goal is to **stack a Pyramid out of any three blocks**, which means the robot need to construct a pyramid-like structure by selecting any three available blocks: two as the base and one stacked on top. The key challenge lies in both object selection and precise spatial configuration. Specifically, the lateral distance between the two base blocks must be carefully chosen to ensure that the top block can be stably placed across them. Furthermore, this task is especially sensitive to execution noise: even if a plan passes all physical constraint checks in simulation, the same stack may collapse in the real environment due to minor perturbations such as control inaccuracy or object pose estimation errors. PRoC3S [3] lacks a feedback mechanism to detect post-simulation failures. Once a plan passes simulation, it is executed directly without verifying whether the real-world outcome satisfies the goal. Our method addresses this by introducing a VLM-based validation step: the executed scene is rendered and checked against the original goal, with replanning initiated if the structure is incorrect.

Task	Symbolic Constraints
Pyramid	<code>StableOn(?block_top, ?block_bottom);</code> <code>ProximitySafe(?block_bottom);</code> <code>PathClear(?gripper, ?block);</code> <code>Before(place(block_top), place(?block_bottom));</code> <code>Distance(?block_bottom) ∈ [0, 0.04]</code> <code>Aligned(?block_bottom) ∧ On(block_top, ?block_bottom);</code> <code>Order(Place(block_bottom1), Place(block_bottom2), Place(block_top));</code> <code>Contact(block_top, ?block_bottom) ∧ CloseTogether(?block_bottom);</code> <code>CenterOfMass(block_top) ∈ SupportArea(?block_bottom)</code>
Line	<code>StableOn(?block, ?table);</code> <code>Alignment(?block);</code> <code>Alignment_tolerance ∈ [-0.01, 0.01]</code>
Packing (Blocks)	<code>ProximitySafe(?block, ?boundary);</code> <code>PlacementFeasible(?block, square_region);</code> <code>WithinBounds(?block, region_center, 2*block_size);</code> <code>ProximitySafe(block, region_center) ∧</code> <code>WithinDistance(?block, region_center, ?max_distance)</code>
Unstack (Blocks)	<code>StableOn(?block, ?bowl);</code> <code>ProximitySafe(?gripper, ?block);</code> <code>ClearOf(?gripper, ?obstacle);</code> <code>PlacementFeasible(green_block, green_bowl);</code> <code>Inside(green_block, green_bowl);</code> <code>Aligned(block_center, bowl_center);</code> <code>NotStacked(green_block, ?obstacle) ∧ offset_z &gt; 0.03</code>
Packing (YCB)	<code>ProximitySafe(?gripper, ?object);</code> <code>GraspFeasible(?grasp, ?object_pose);</code> <code>ProximitySafe(object, table_center);</code> <code>Distance(?object, table_center) &lt; 0.06;</code> <code>PlacementFeasible(?object, center, threshold) ∧ threshold = 0.06;</code> <code>Graspable(?grasp, ?object);</code> <code>CollisionFree(?gripper, ?object)</code>
Stacking (YCB)	<code>OnTop(object_a, object_b);</code> <code>ClearOf(object_a, surface);</code> <code>ClearSurface(object_b);</code> <code>PlacementFeasible(object_a, object_b);</code> <code>StableOn(object_a, object_b);</code> <code>AlignedForStacking(object_a, object_b);</code> <code>Graspable(?object, ?grasp) ∧ CollisionFree(?object)</code>

Table 4: Inducted constraints for each task across Arrange-Blocks and Arrange-YCB domains

612 In the first round of code generation, InstructFlow receives the goal: “*Stack a pyramid out of any*  
613 *three blocks.*” The **InstructFlow Planner** constructs an initial instruction graph consisting of seven  
614 planning nodes, each representing a subgoal in the pyramid assembly process. Guided by this graph,  
615 the code generator produces an executable plan. However, after executing the plan, VLM-based  
616 semantic validation reports a failure, indicating the following issue:

[Error Message]: "The blocks are not arranged in a pyramid formation"

617 Given the failure feedback and generated code, the **Constraint Generator** localizes the issue to the  
618 code block responsible for pyramid construction, specifically the subgoal corresponding to subgoal15  
619 in `get_plan()` and the `offset_x` parameter defined in `get_domain()`. By examining the sampled  
620 parameter values, InstructFlow infers that the current range of `offset_x` is too wide, causing the  
621 base blocks to be placed too far apart. As a result, the top block either falls during execution or  
622 forms a configuration that is not recognized as a pyramid by the VLM. This reasoning leads to the  
623 generation of an explicit symbolic constraint:



```

1  def gen_plan(initial: RavenBelief,
2      base_x: float, base_y: float,
3      offset_x: float):
4      plan = []
5
6      # Subgoal 1: Select three blocks for the pyramid
7      blocks = [obj for obj in initial.objects.values() if obj.category == "block"]
8      block1, block2, block3 = blocks[:3]
9
10     # Subgoal 2: Pick the first block for the base of the pyramid
11     plan.append(Action("pick", block1.pose.point))
12
13     # Subgoal 3: Place the first block at the base position
14     plan.append(Action("place", [base_x, base_y, block1.pose.point[2]]))
15
16     # Subgoal 4: Pick the second block for the base of the pyramid
17     plan.append(Action("pick", block2.pose.point))
18
19     # Subgoal 5: Place the second block next to the first block at the base position, ensuring alignment
20     plan.append(Action("place", [base_x + offset_x, base_y, block2.pose.point[2]]))
21
22     # Subgoal 6: Pick the third block for the top of the pyramid
23     plan.append(Action("pick", block3.pose.point))
24
25     # Subgoal 7: Place the third block on top of the first and second blocks to complete the pyramid
26     plan.append(Action("place", [base_x + offset_x / 2, base_y, block1.pose.point[2] + BLOCK_SIZE]))
27
28     return plan
29
30 def gen_domain(initial: RavenBelief):
31     return {
32         "base_x": ContinuousSampler(TABLE_BOUNDS[0][0], TABLE_BOUNDS[0][1] - BLOCK_SIZE),
33         "base_y": ContinuousSampler(TABLE_BOUNDS[1][0], TABLE_BOUNDS[1][1]),
34         "offset_x": ContinuousSampler(BLOCK_SIZE, BLOCK_SIZE * 2),
35         "offset_y": ContinuousSampler(0, BLOCK_SIZE),
36     }

```

Figure 8: A code snippet illustrating how InstructFlow repairs the **Pyramid** task plan: to correct the code, **InstructFlow** incorporates a **parameter range refinement node** into the instruction graph based on the inducted symbolic constraint, which guides the regenerated code to reduce the sampling range of **offset\_x** in the **get\_domain** function. This adjustment reduces the spacing between base blocks, enabling a **stable top placement** and successful pyramid construction.

$$\phi_{\text{dist}} := \text{Distance}(\text{?block\_bottom}) \in [0, 0.04]$$

624 Here, 0.04 corresponds to the environment-defined block size, ensuring that the top block can span  
625 both base blocks without falling or misalignment.

626 Given the symbolic constraint, **InstructFlow Planner** dynamically updates the instruction graph by  
627 introducing reasoning nodes that refine the value range of **offset\_x**. Specifically, the updated graph  
628 incorporates a parameter range refinement node  $v_{\text{param}}^{\text{reason}}$ , whose output is an instruction to  
629 narrow the sampling bounds of **offset\_x**.

630 Based on the updated instruction graph, the **Code Generator** regenerates the executable code  
631 to satisfy the induced symbolic constraints. As shown in Figure 5, the regenerated code mod-  
632 ifies the **get\_domain()** function by narrowing the sampling range of **offset\_x** from the original  
633 (**BLOCK\_SIZE**, **BLOCK\_SIZE \* 2**) to (**0**, **BLOCK\_SIZE**). This adjustment results in a tighter ar-  
634 rangement of the base blocks, enabling the top block to rest stably and form a recognizable pyramid  
635 structure. The final scene passes VLM validation, and the task is successfully completed.

636 **Stacking (Arrange-YCB).** This task requires the agent to stack one object on top of another,  
637 selected from a diverse set of YCB objects with varied shapes, sizes, and physical properties.  
638 The primary challenge lies in object selection: due to irregular geometries and asymmetric mass  
639 distributions, not all object pairs are feasible for stable stacking. Therefore, selecting an appropriate  
640 pair is critical to the success of the task. A second challenge arises from the lack of feedback handling  
641 in PRoC3S. Due to shape mismatch or small execution disturbances, the stacked configuration may  
642 fail in the real environment but without replanning.

643 In the first round of code generation, InstructFlow receives the goal: "Stack any object on any  
644 other object". The **InstructFlow Planner** constructs an initial instruction graph consisting of  
645 three planning nodes, corresponding to selecting two objects, picking the first object, and placing it

```

1  def gen_plan(initial: RavenBelief, object_name1: str, grasp1: RavenGrasp,
2  -   object_name2: str, grasp2: RavenGrasp):
3  +   def gen_plan(initial: RavenBelief, grasp: RavenGrasp):
4
5      plan = []
6
7      #Subgoal 1: Pick the first selected object
8      plan.append(Action("pick", [object_name1, grasp1]))
9
10     #Subgoal 2: Get the pose of the second object to stack on top
11     second_object_pose = initial.objects[object_name2].pose
12
13     #Subgoal 3: Place the first object on top of the second object
14     plan.append(Action("place", [object_name1, grasp1, RavenPose(x=second_object_pose.x,
15     y=second_object_pose.y, z=second_object_pose.z + BLOCK_SIZE)]))
16
17     #Subgoal 1: Define the objects to be used in the plan
18     object_name1 = "object_2" # apple
19     object_name2 = "object_4" # power_drill
20
21     #Subgoal 2: Pick the first selected object (apple)
22     plan.append(Action("pick", [object_name1, grasp]))
23
24     #Subgoal 3: Get the pose of the second object (power_drill) to stack on top
25     second_object_pose = initial.objects[object_name2].pose
26
27     #Subgoal 4: Place the first object (apple) on top of the second object (power_drill)
28     plan.append(Action("place", [object_name1, grasp, RavenPose(x=second_object_pose.x,
29     y=second_object_pose.y, z=second_object_pose.z + BLOCK_SIZE)]))
30
31     return plan
32
33     def gen_domain(initial: RavenBelief):
34     object_ids = list(initial.objects.keys())
35     return {
36         "object_name1": DiscreteSampler(object_ids),
37         "grasp1": GraspSampler(),
38         "object_name2": DiscreteSampler(object_ids),
39         "grasp2": GraspSampler(),
40     }
41
42     +   return {
43         "grasp": GraspSampler(),
44     }
45
46

```

Figure 9: A code snippet illustrating how **InstructFlow** repairs the **Stacking** task plan: to correct the code, **InstructFlow** updates the instruction graph based on the induced symbolic constraint. This update adjusts the subgoals from **randomly selecting stacking targets** to using a **specifically determined object pair**, guiding the regenerated code to modify the corresponding logic in **get\_plan** and remove the random sampling from **get\_domain**. This adjustment ensures the **physical feasibility** of the stacking operation.

on top of the second. Guided by this structure, the **Code Generator** generates an executable plan. However, after execution, VLM-based semantic validation reports a failure, indicating that the chosen object pair did not result in a valid or stable stacked configuration.

[Error Message]: "No objects are stacked on top of each other"

Given the failure feedback and generated code, the **Constraint Generator** localizes the issue to the code block responsible for selecting object pairs. Specifically, the analysis reveals that the object pair was selected via an environment-defined `DiscreteSampler()`, which randomly samples two objects without considering any factors that influence stackability. As a result, the constraint generator induces the following symbolic constraint to guide future selection:

$$\phi_{\text{stack}} := \text{AlignedForStacking}(\text{object\_a}, \text{object\_b}) \wedge \text{PlacementFeasible}(\text{object\_a}, \text{object\_b})$$

Given the symbolic constraint, **InstructFlow Planner** dynamically updates the instruction graph by introducing reasoning nodes that guide the selection of a physically compatible object pair for stacking. Specifically, the updated graph includes a `spatial_relation_reasoning_node`  $v_{\mathcal{T}^{\text{spatial}}}^{\text{reason}}$  to analyze relative object positions and an `object_density_analysis_node`  $v_{\mathcal{T}^{\text{density}}}^{\text{reason}}$  to evaluate the local clutter surrounding each object. These outputs are fed into an `object_selection_reasoning_node`  $v_{\mathcal{T}^{\text{select}}}^{\text{reason}}$ , which selects `object_2` and `object_4`

660 as a feasible stacking pair based on the symbolic constraint `SelectStackablePair`. A  
 661 `plan_logic_reasoning_node`  $v_{\mathcal{T}_{plan}}^{\text{reason}}$  then determines the appropriate execution order, generating  
 662 a subgoal sequence: “Pick object\_2” followed by “Place object\_2 on top of object\_4.”  
 663 Based on the updated instruction graph, the **Code Generator** regenerates the executable code to  
 664 satisfy the induced symbolic constraints. As shown in Figure 6, the regenerated code eliminates the  
 665 use of unconstrained object sampling in `gen_domain()`, replacing it with fixed object assignments  
 666 (object\_2 and object\_4) determined through reasoning nodes. This ensures that the selected  
 667 object pair adheres to the `AlignedForStacking` and `PlacementFeasible` constraints. The new  
 668 plan simplifies the domain by sampling only the grasp for the manipulated object and directly  
 669 encodes the stacking intent in the plan logic. With this targeted refinement, the resulting scene passes  
 670 VLM-based check, and the stacking task is successfully completed.