

466 **A Environments**

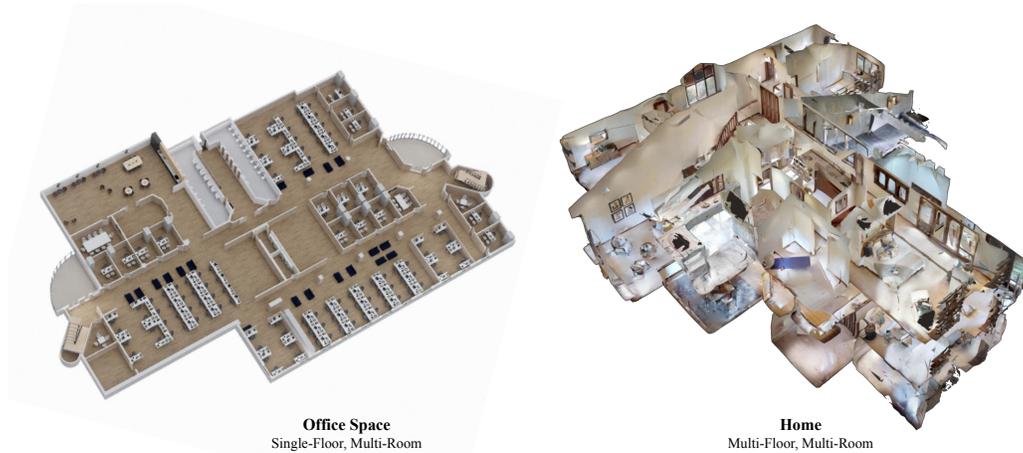


Figure 3: **Large-scale environments used to evaluate SayPlan.**

467 We evaluate SayPlan across a set of two large-scale environments spanning multiple rooms and
468 floors as shown in Figure 3. We provide details of each of these environments below:

469 **Office:** A large-scale office floor, spanning 36 rooms and 150 assets and objects which the agent
470 can interact with. This scene graph represents a real-world office floor within which a mobile ma-
471 nipulator robot is present. This allows us to embody the plans generated using SayPlan and evaluate
472 their feasibility in the corresponding environment. A full visual 3D scene graph representation of
473 the office floor is provided in Appendix C and D. We utilise generic names across our scene to
474 maintain anonymity throughout the review process.

475

476 **Home:** An existing 3D scene graph from the Stanford 3D Scene Graph dataset [11] which consists
477 of a family home environment (*Klickitat*) spanning 32 rooms across 3 floors and contains 121
478 assets and objects which the agent can interact with. A 3D visual of this environment can be viewed
479 at the 3D Scene Graph project website.²

480 **B Tasks**

481 We evaluate SayPlan across 4 instruction sets which are classified to evaluate different aspects of its
482 3D scene graph reasoning and planning capabilities:

483 **Simple Search:** Focused on evaluating the semantic search capabilities of the LLM based on queries
484 which directly reference information in the scene graph as well as the basic graph-based reasoning
485 capabilities of the LMM.

486 **Complex Search:** Abstract semantic search queries which require complex reasoning. The infor-
487 mation required to solve these search tasks is not readily available in the graph and has to be inferred
488 by the underlying LLM.

489 **Simple Planning:** Task planning queries which require the agent to perform graph search, causal
490 reasoning and environment interaction in order to solve the task. Typically requires shorter horizon
491 plans over single rooms.

492 **Long Horizon Planning:** Long Horizon planning queries require multiple interactive steps. These
493 queries evaluate SayPlan’s ability to reason over temporally extended instructions to investigate how
494 well it scales to such regimes. Typically requires long horizon plans spanning multiple rooms.

495 The full list of instructions used and the corresponding aspect the query evaluates are given in the
496 following tables.

²3dscenegraph.stanford.edu/Klickitat

497 **B.1 Simple Search**

498 **B.1.1 Office Environment**

Instruction	
Find me object K31X.	▷ unguided search with no semantic cue
Find me a carrot.	▷ semantic search based on node name
Find me anything purple in the postdoc bays.	▷ semantic search with termination conditioned on attribute
Find me a ripe banana.	▷ semantic search with termination conditioned on attribute
Find me something that has a screwdriver in it.	▷ unguided search with termination conditioned on children
One of the offices has a poster of the Terminator. Which one is it?	▷ semantic search with termination conditioned on children
I printed a document but I don't know which printer has it. Find the document.	▷ semantic search based on parent
I left my headphones in one of the meeting rooms. Locate them.	▷ semantic search based on parent
Find the PhD bay that has a drone in it.	▷ semantic search with termination conditioned on children
Find the kale that is not in the kitchen.	▷ semantic search with termination conditioned on a negation predicate on parent
Find me an office that does not have a cabinet.	▷ semantic search with termination conditioned on a negation predicate on children
Find me an office that contains a cabinet, a desk, and a chair.	▷ semantic search with termination conditioned on a conjunctive query on children
Find a book that was left next to a robotic gripper.	▷ semantic search with termination conditioned on a sibling
Luis gave one of his neighbours a stapler. Find the stapler.	▷ semantic search with termination conditioned on a sibling
There is a meeting room with a chair but no table. Locate it.	▷ semantic search with termination conditioned on a conjunctive query with negation

Table 5: **Simple Search Instructions.** Evaluated in Office Environment.

Instruction	
Find me a FooBar.	▷ unguided search with no semantic cue
Find me a bottle of wine.	▷ semantic search based on node name
Find me a plant with thorns.	▷ semantic search with termination conditioned on attribute
Find me a plant that needs watering.	▷ semantic search with termination conditioned on attribute
Find me a bathroom with no toilet.	▷ semantic search with termination conditioned on a negation predicate
The baby dropped their rattle in one of the rooms. Locate it.	▷ semantic search based on node name
I left my suitcase either in the bedroom or the living room. Which room is it in.	▷ semantic search based on node name
Find the room with a ball in it.	▷ semantic search based on node name
I forgot my book on a bed. Locate it.	▷ semantic search based on node name
Find an empty vase that was left next to sink.	▷ semantic search with termination conditioned on sibling
Locate the dining room which has a table, chair and a baby monitor.	▷ semantic search with termination conditioned on conjunctive query
Locate a chair that is not in any dining room.	▷ semantic search with termination conditioned on negation predicate
I need to shave. Which room has both a razor and shaving cream.	▷ semantic search with termination conditioned on children
Find me 2 bedrooms with pillows in them.	▷ semantic search with multiple returns
Find me 2 bedrooms without pillows in them.	▷ semantic search with multiple returns based on negation predicate

Table 6: **Simple Search Instructions.** Evaluated in Home Environment.

500 **B.2 Complex Search**

501 **B.2.1 Office Environment**

Instruction	
Find object J64M. J64M should be kept at below 0 degrees Celsius.	▷ semantic search guided by implicit world knowledge (knowledge not directly encoded in graph)
Find me something non vegetarian.	▷ semantic search with termination conditioned on implicit world knowledge
Locate something sharp.	▷ unguided search with termination conditioned on implicit world knowledge
Find the room where people are playing board games.	▷ semantic search with termination conditioned on ability to deduce context from node children using world knowledge (“board game” is not part of any node name or attribute in this graph)
Find an office of someone who is clearly a fan of Arnold Schwarzenegger.	▷ semantic search with termination conditioned on ability to deduce context from node children using world knowledge
There is a postdoc that has a pet Husky. Find the desk that’s most likely theirs.	▷ semantic search with termination conditioned on ability to deduce context from node children using world knowledge
One of the PhD students was given more than one complimentary T-shirts. Find his desk.	▷ semantic search with termination conditioned on the number of children
Find me the office where a paper attachment device is inside an asset that is open.	▷ semantic search with termination conditioned on node descendants and their attributes
There is an office which has a cabinet containing exactly 3 items in it. Locate the office.	▷ semantic search with termination conditioned on the number of children
There is an office which has a cabinet containing a rotten apple. The cabinet name contains an even number. Locate the office.	▷ semantic search guided by numerical properties
Look for a carrot. The carrot is likely to be in a meeting room but I’m not sure.	▷ semantic search guided by used provided bias
Find me a meeting room with a RealSense camera.	▷ semantic search that has no result (no meeting room has a realsense camera in the graph)
Find the closest fire extinguisher to the manipulation lab.	▷ search guided by node distance
Find me the closest meeting room to the kitchen.	▷ search guided by node distance
Either Filipe or Tobi has my headphones. Locate it.	▷ evaluating constrained search, early termination once the two office are explored

Table 7: **Complex Search Instructions.** Evaluated in Office Environment.

Instruction	
I need something to access ChatGPT. Where should I go?	▷ semantic search guided by implicit world knowledge
Find the livingroom that contains the most electronic devices.	▷ semantic search with termination conditioned on children with indirect information
Find me something to eat with a lot of potassium.	▷ semantic search with termination conditioned on implicit world knowledge
I left a sock in a bedroom and one in the living room. Locate them. They should match.	▷ semantic search with multiple returns
Find me a potted plant that is most likely a cactus.	▷ semantic search with termination implicitly conditioned on attribute
Find the dining room with exactly 5 chairs.	▷ semantic search with termination implicitly conditioned on quantity of children
Find me the bedroom closest to the home office.	▷ semantic search with termination implicitly conditioned on node distance
Find me a bedroom with an unusual amount of bowls.	▷ semantic search with termination implicitly conditioned on quantity of children
Which bedroom is empty.	▷ semantic search with termination implicitly conditioned on quantity of children
Which bathroom has the most potted plants.	▷ semantic search with termination implicitly conditioned on quantity of children
The kitchen is flooded. Find somewhere I can heat up my food.	▷ semantic search guided by negation
Find me the room which most likely belongs to a child	▷ semantic search with termination conditioned on ability to deduce context from node children using world knowledge
15 guests are arriving. Locate enough chairs to seat them.	▷ semantic search with multiple returns
A vegetarian dinner was prepared in one of the dining rooms. Locate it.	▷ semantic search with selection criteria based on world knowledge
My tie is in one of the closets. Locate it.	▷ evaluating constrained search, termination after exploring closets

Table 8: **Complex Search Instructions.** Evaluated in Home Environment.

503 **B.3 Simple Planning**

Instruction
Close Jason's cabinet.
Refrigerate the orange left on the kitchen bench.
Take care of the dirty plate in the lunchroom.
Place the printed document on Will's desk.
Peter is working hard at his desk. Get him a healthy snack.
Hide one of Peter's valuable belongings.
Wipe the dusty admin shelf.
There is coffee dripping on the floor. Stop it.
Place Will's drone on his desk.
Move the monitor from Jason's office to Filipe's.
My parcel just got delivered! Locate it and place it in the appropriate lab.
Check if the coffee machine is working.
Heat up the chicken kebab.
Something is smelling in the kitchen. Dispose of it.
Throw what the agent is holding in the bin.

Table 9: **Simple Planning Instructions.** Evaluated in Office Environment.

504 **B.4 Long Horizon Planning**

Instruction
Heat up the noodles in the fridge, and place it somewhere where I can enjoy it.
Throw the rotting fruit in Dimity's office in the correct bin.
Wash all the dishes on the lunch table. Once finished, place all the clean cutlery in the drawer.
Safely file away the freshly printed document in Will's office then place the undergraduate thesis on his desk.
Make Niko a coffee and place the mug on his desk.
Someone has thrown items in the wrong bins. Correct this.
Tobi spilt soda on his desk. Throw away the can and take him something to clean with.
I want to make a sandwich. Place all the ingredients on the lunch table.
A delegation of project partners is arriving soon. We want to serve them snacks and non-alcoholic drinks. Prepare everything in the largest meeting room. Use items found in the supplies room only.
Serve bottled water to the attendees who are seated in meeting room 1. Each attendee can only receive a single bottle of water.
Empty the dishwasher. Place all items in their correct locations
Locate all 6 complimentary t-shirts given to the PhD students and place them on the shelf in admin.
I'm hungry. Bring me an apple from Peter and a pepsi from Tobi. I'm at the lunch table.
Let's play a prank on Niko. Dimity might have something.
There is an office which has a cabinet containing a rotten apple. The cabinet name contains an even number. Locate the office, throw away the fruit and get them a fresh apple.

Table 10: **Long-Horizon Planning Instructions.** Evaluated in Office Environment.



Figure 4: 3D Scene Graph - Fully Expanded Office Environment. Full 3D scene graph exposing all the rooms, assets and objects available in the scene. Note that the LLM agent never sees all this information unless it chooses to expand every possible node without contraction.

506 **D Contracted 3D Scene Graph: Office Environment**



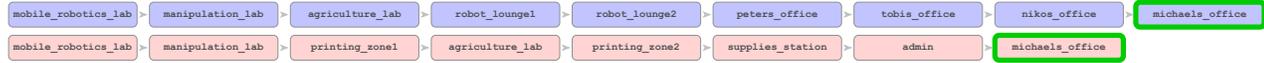
Figure 5: **3D Scene Graph - Contracted Office Environment.** Contracted 3D scene graph exposing only the highest level within the hierarchy - room nodes. This results in an 82.1% reduction in the number of tokens required to represent the scene before the exploration phase.

507 **E Semantic Search Evaluation Results**

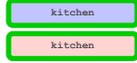
508 - Full listings of the generated semantic search sequences for the evaluation instruction sets are
509 provided on the following page -



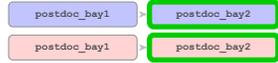
Find me object K31X.



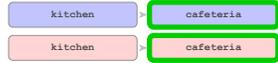
Find me a carrot.



Find me anything purple in the postdoc bays.



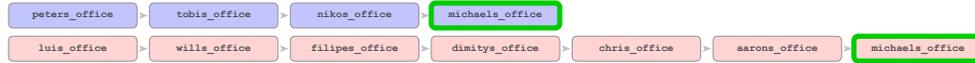
Find me a ripe banana.



Find me something that has a screwdriver in it.



One of the offices has a poster of the Terminator. Which one is it?



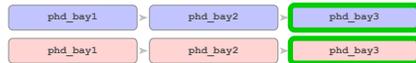
I printed a document, but I dont know which printer has it. Find the document.



I left my headphones in one of the meeting rooms. Locate them.



Find the PhD bay that has a drone in it.



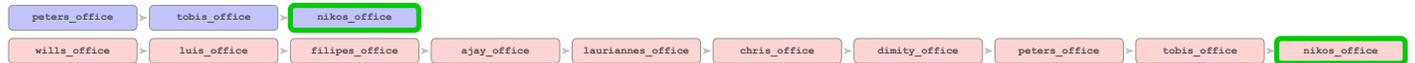
Find the kale that is not in the kitchen.



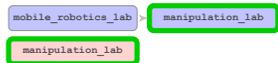
Find me an office that does not have a cabinet.



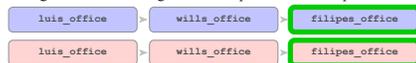
Find me an office that contains a cabinet, a desk and a chair.



Find me a book that was left next to a robotic gripper.



Luis gave one of his neighbours a stapler. Find the stapler.



There is a meeting room with a chair but no table. Locate it.

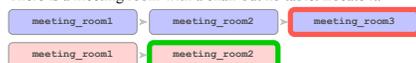
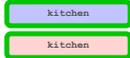


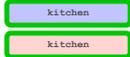
Table 11: **Simple Search Office Environment Evaluation.** Sequence of Explored Nodes for Simple Search Office Environment Instructions.



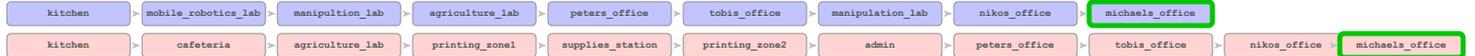
Find object J64M. J64M should be kept at below 0 degrees Celsius.



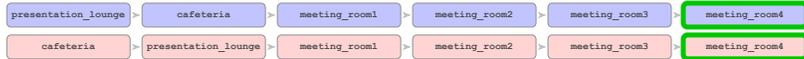
Find me something non vegetarian.



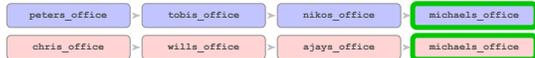
Locate something sharp.



Find the room where people are playing board games.



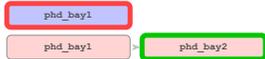
Find the office of someone who is clearly a fan of Arnold Schwarzenegger.



There is postdoc that has a pet Husky. Find the desk that's most likely theirs.



One of the PhD students was given more than one complimentary T-shirt. Find his desk.



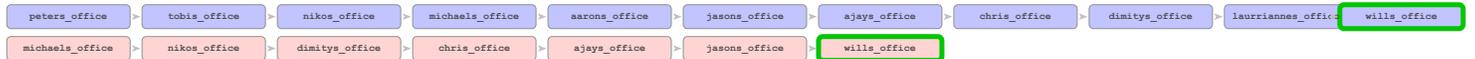
Find me the office where a paper attachment device is inside an asset that is open.



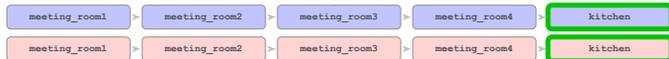
There is an office which has a cabinet containing exactly 3 items in it. Locate the office.



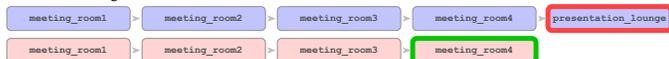
There is an office containing a rotten apple. The cabinet name contains an even number. Locate the office.



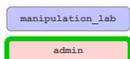
Look for a carrot. The carrot is likely to be in a meeting room but I'm not sure.



Find me a meeting room with a RealSense camera.



Find the closest fire extinguisher to the manipulation lab.



Find me the closest meeting room to the kitchen.



Either Filipe or Tobi has my headphones. Locate them.

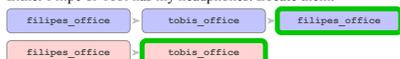
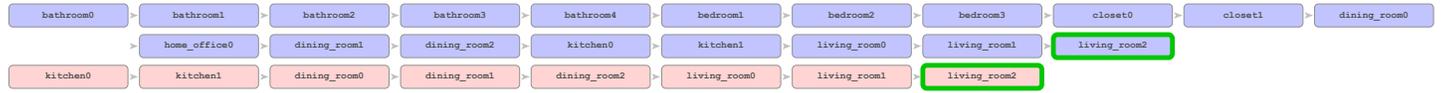


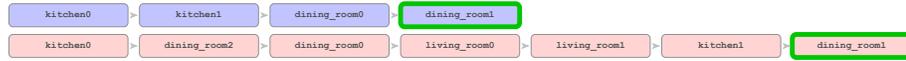
Table 12: **Complex Search Office Environment Evaluation.** Sequence of Explored Nodes for Complex Search Office Environment Instructions.



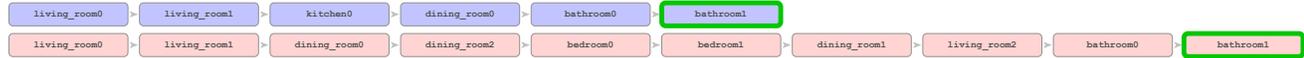
Find me a FooBar.



Find me a bottle of wine.



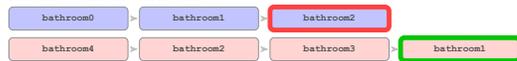
Find me a plant with thorns.



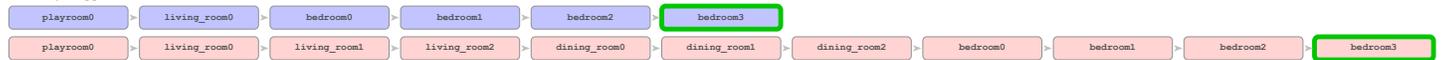
Find me a plant that needs watering.



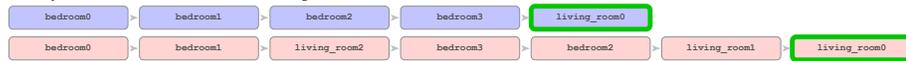
Find me a bathroom with no toilet.



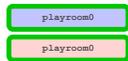
The baby dropped their rattle in one of the rooms. Locate it.



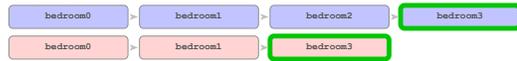
I left my suitcase either in the bedroom or the living room. Which room is it in.



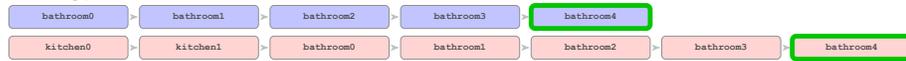
Find the room with a ball in it.



I forgot my book on a bed. Locate it.



Find an empty vase that was left next to a sink.



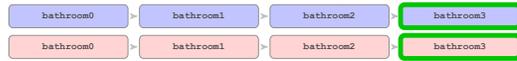
Locate the dining room which has a table, chair and a baby monitor.



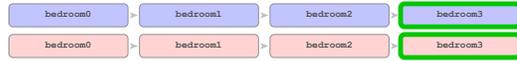
Locate a chair that is not in any dining room.



I need to shave. Which room has both a razor and shaving cream.



Find me 2 bedrooms with pillows in them.



Find me 2 bedrooms without pillows in them.

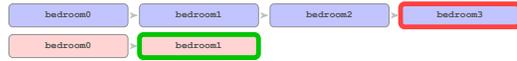
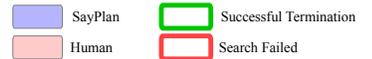
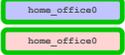


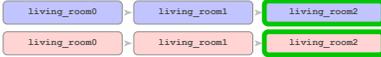
Table 13: **Simple Search Home Environment Evaluation.** Sequence of Explored Nodes for Simple Search Home Environment Instructions.



I need something to access ChatGPT. Where should I go?.



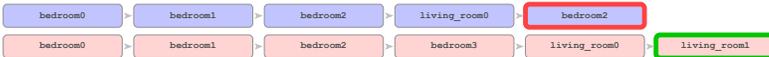
Find the livingroom that contains the most electronic devices.



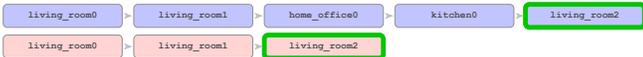
Find me something to eat with alot of potassium.



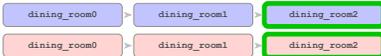
I left a sock in a bedroom and in one of the livingrooms. Locate them. They should match.



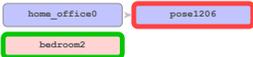
Find the potted plant that is most likely a cactus.



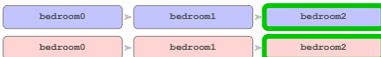
Find the dining room with exactly 5 chairs.



Find me the bedroom closest to the home office.



Find me the bedroom with an unusual amount of bowls.



Which bedroom is empty.



Which bathroom has the most potted plants.



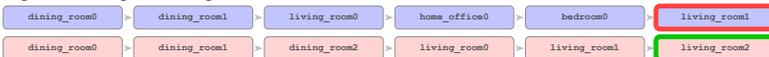
The kitchen is flooded. Find somewhere I can heat up my food.



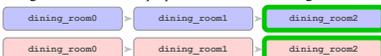
Find me the room which most likley belongs to a child.



15 guests are arriving. Locate enough chairs to seat them.



A vegetarian dinner was prepared in one of the dining rooms. Locate it.



My tie is in one of the closets. Locate it.



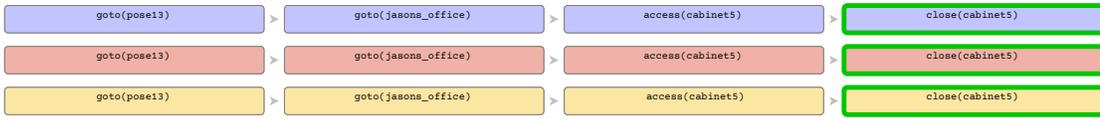
Table 14: **Complex Search Home Environment Evaluation.** Sequence of Explored Nodes for Complex Search Home Environment Instructions.

510 **F Causal Planning Evaluation Results**

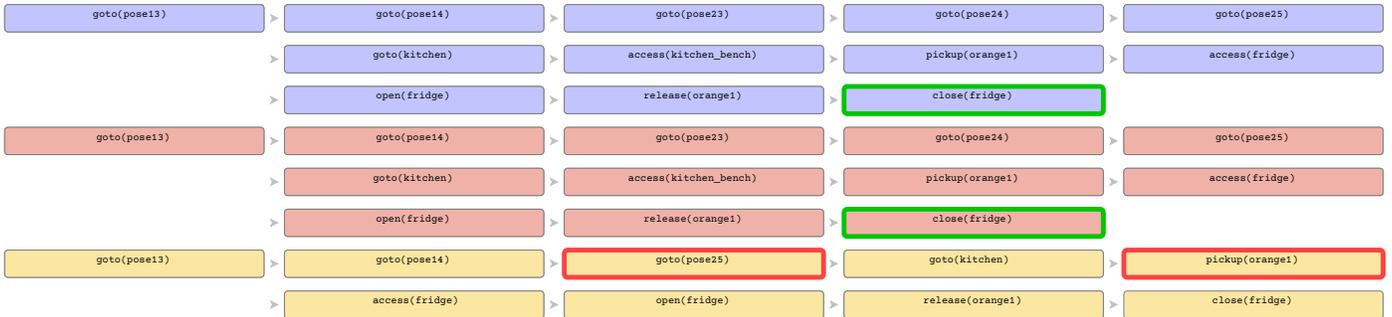
511 - Full listings of the generated planning sequences for the evaluation instruction sets are provided
512 on the following page -



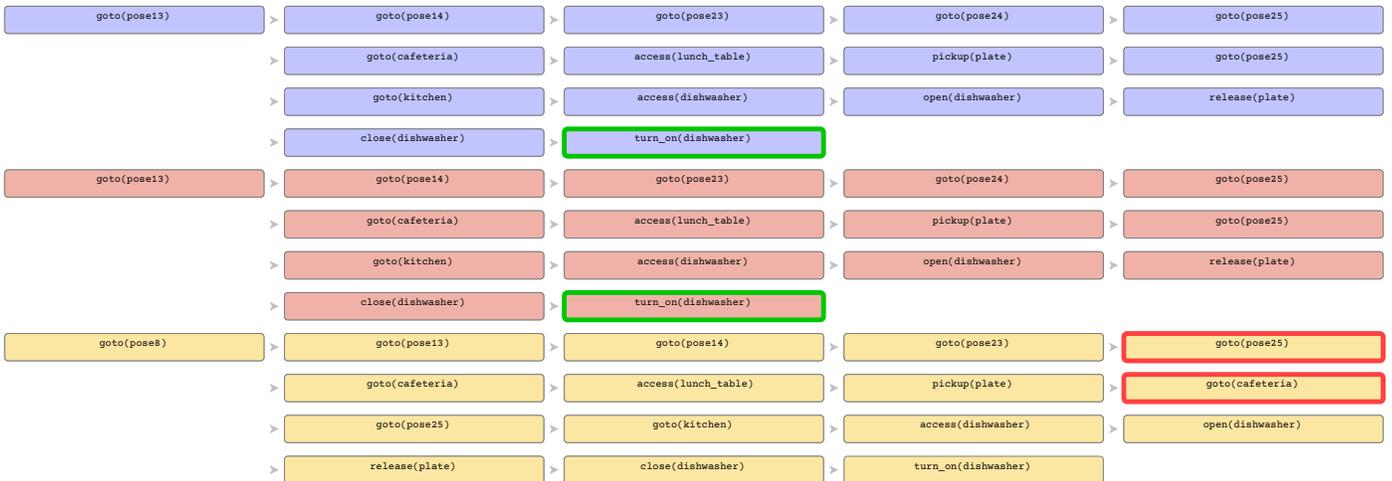
Close Jason's cabinet.



Refrigerate the orange left on the kitchen bench.



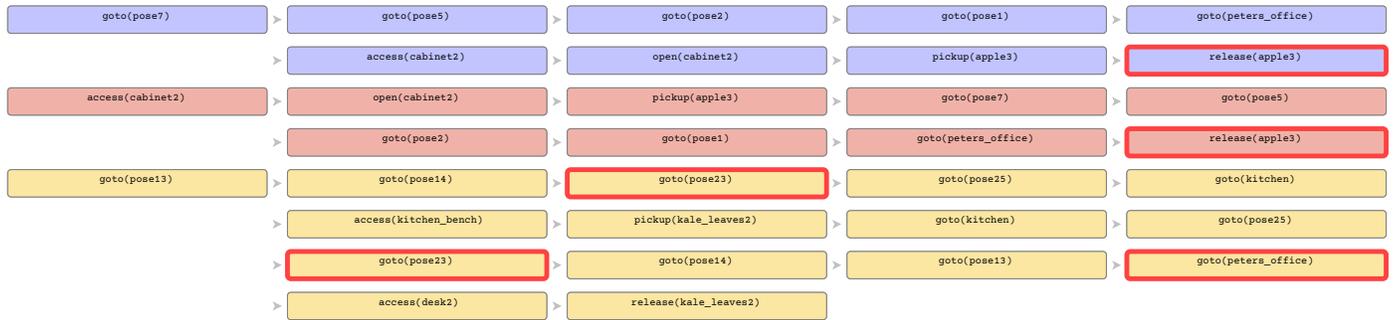
Take care of the dirty plate in the lunchroom.



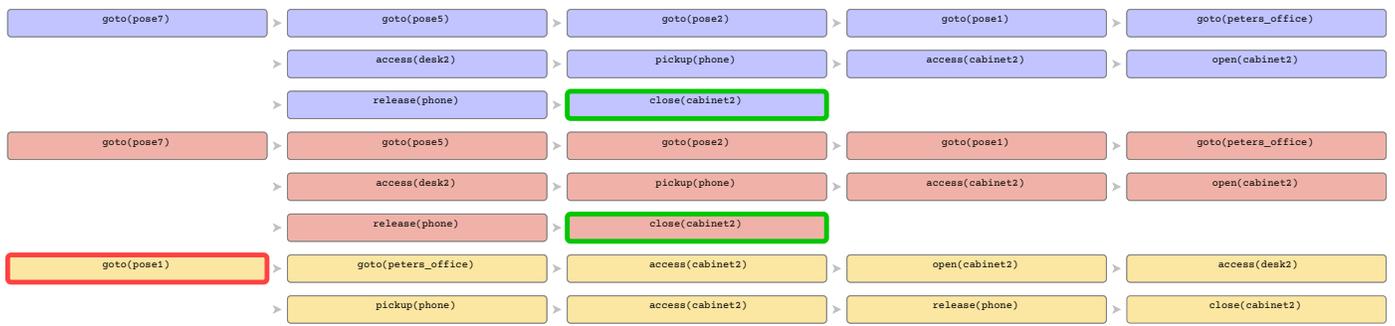
Place the printed document on Will's desk.



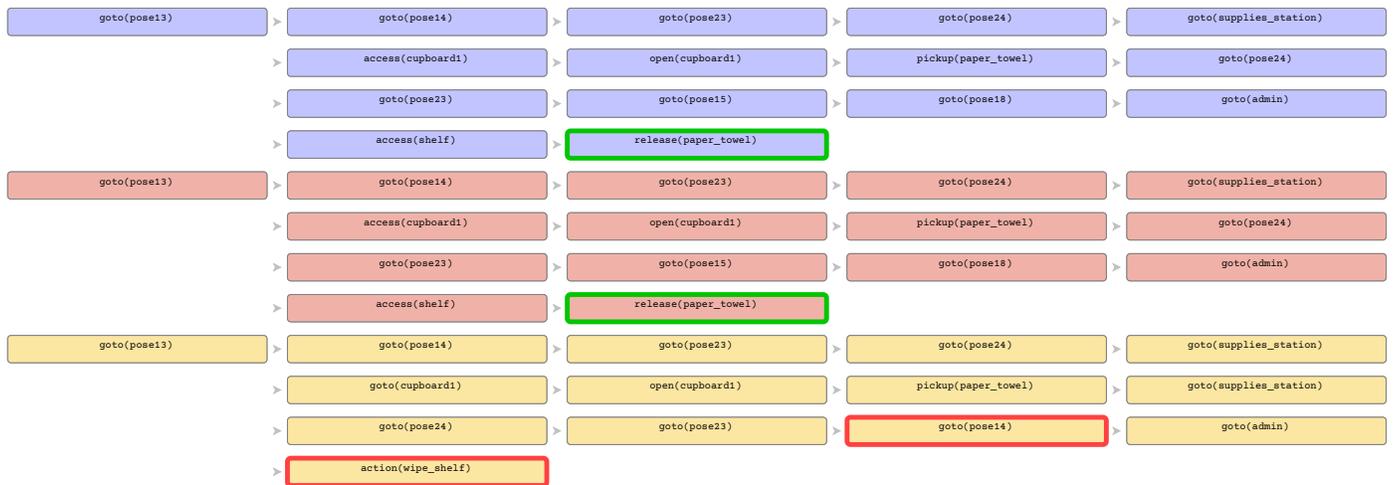
Peter is working hard at his desk. Get him a healthy snack.



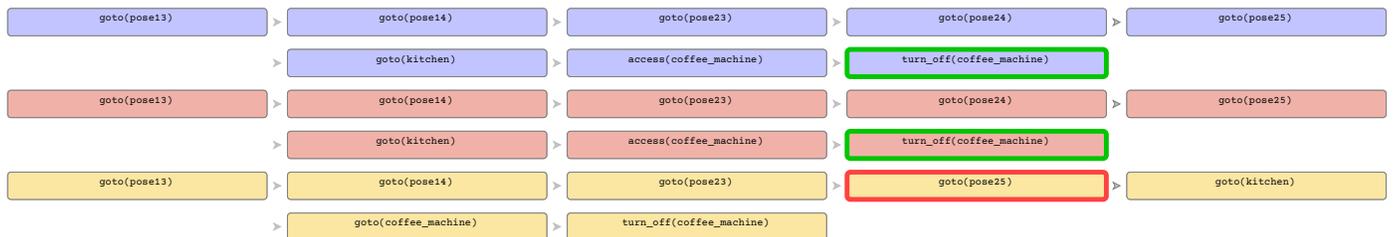
Hide one of Peter's valuable belongings.



Wipe the dusty admin shelf.

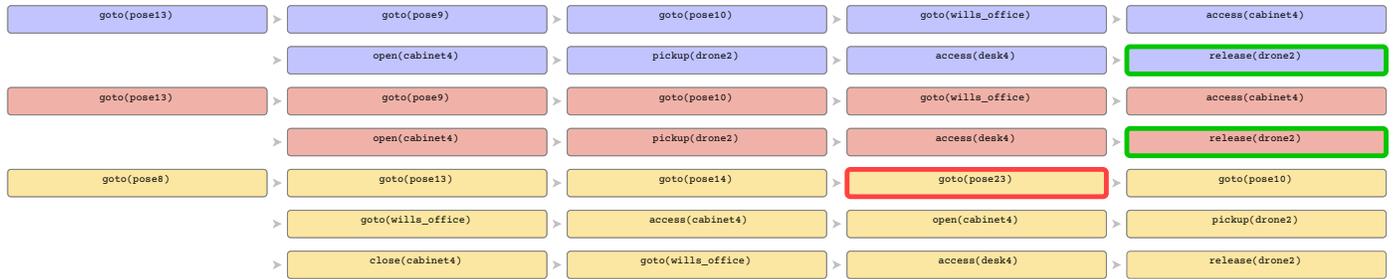


There is coffee dripping on the floor. Stop it.

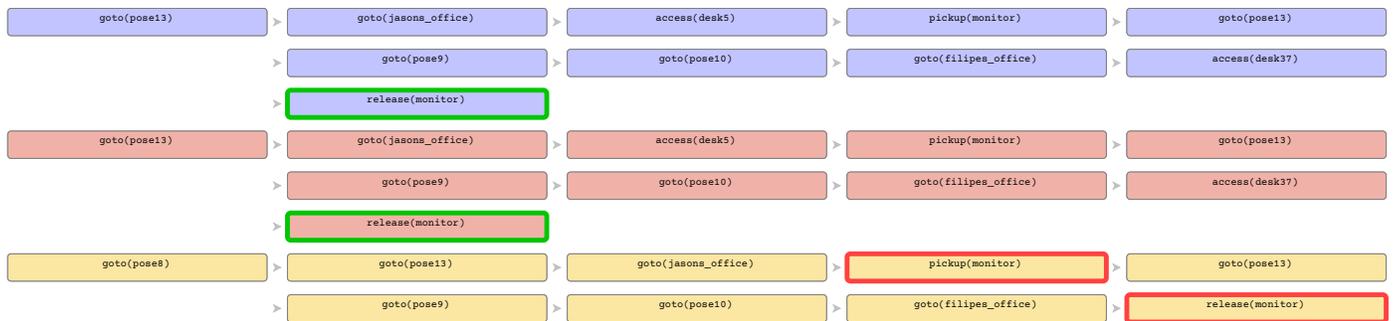




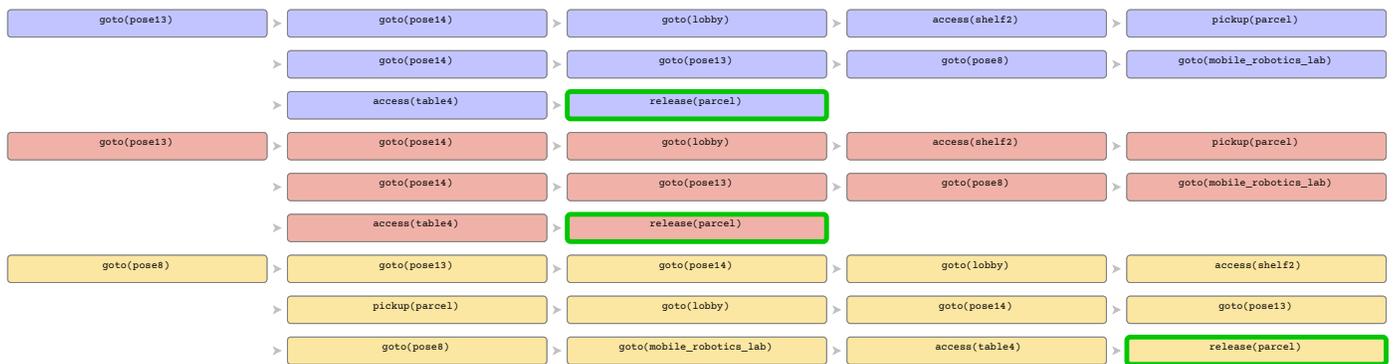
Place Will's drone on his desk.



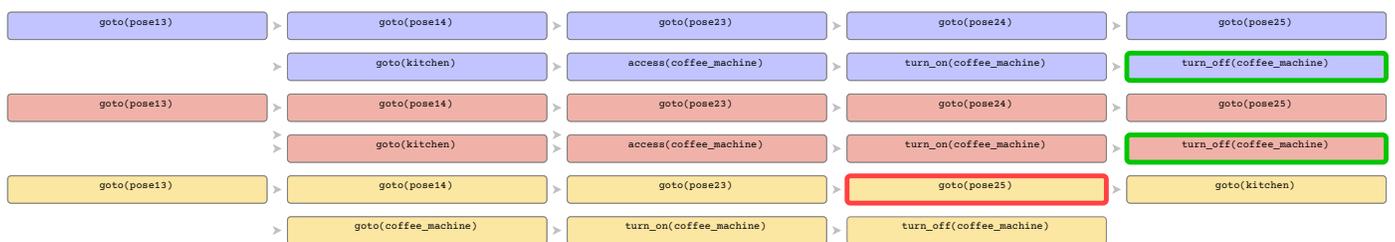
Move the monitor from Jason's office to Filipe's.



My parcel just got delivered! Locate it and place it in the appropriate lab.

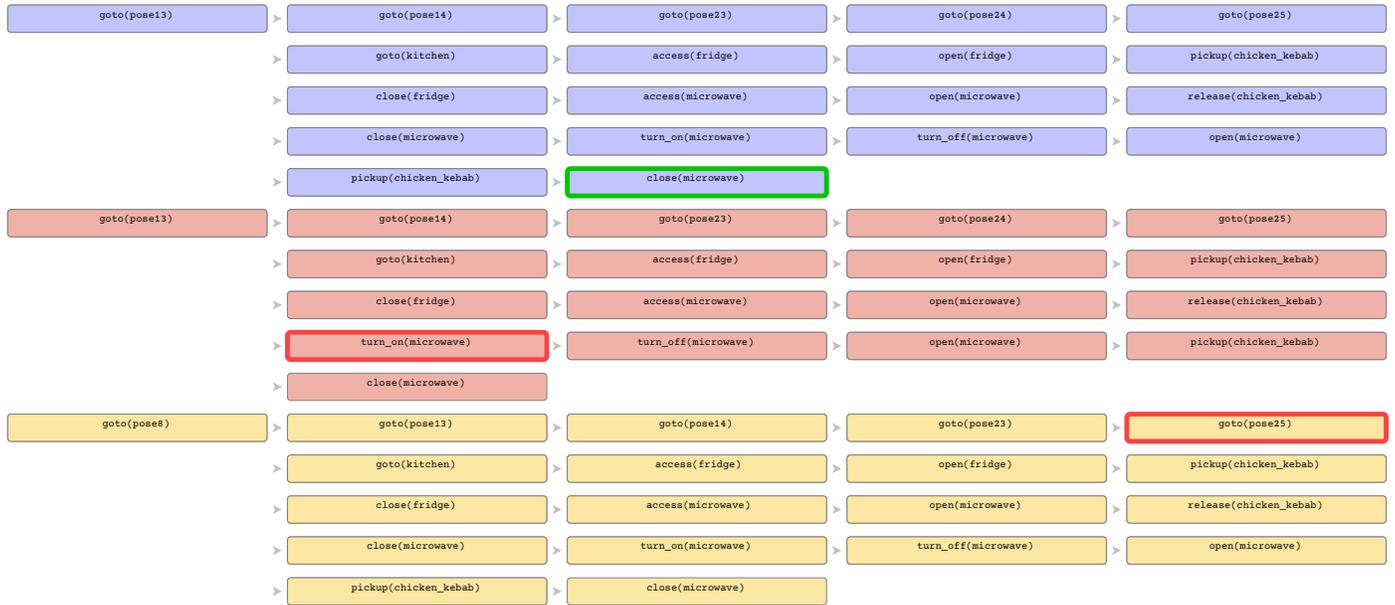


Check if the coffee machine is working.

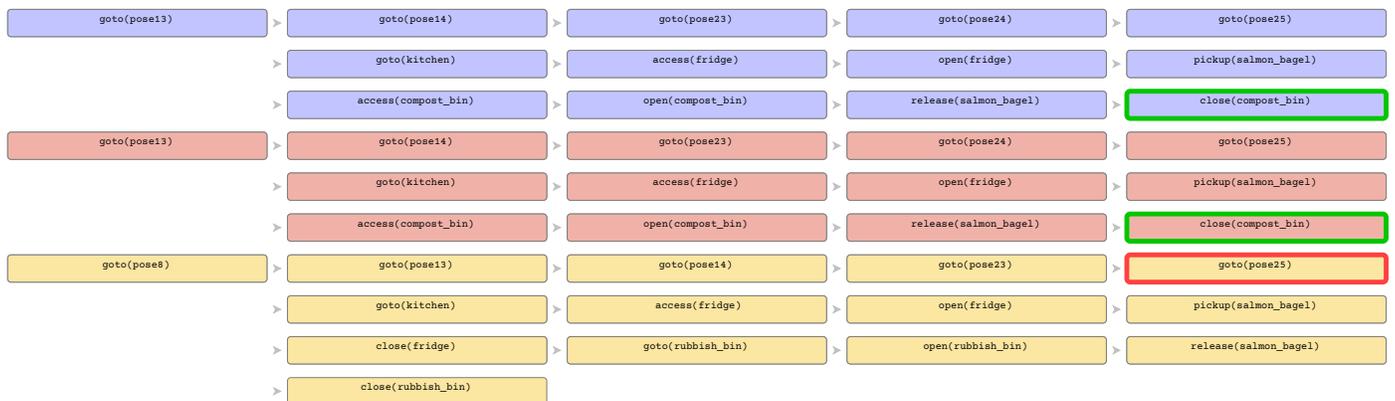




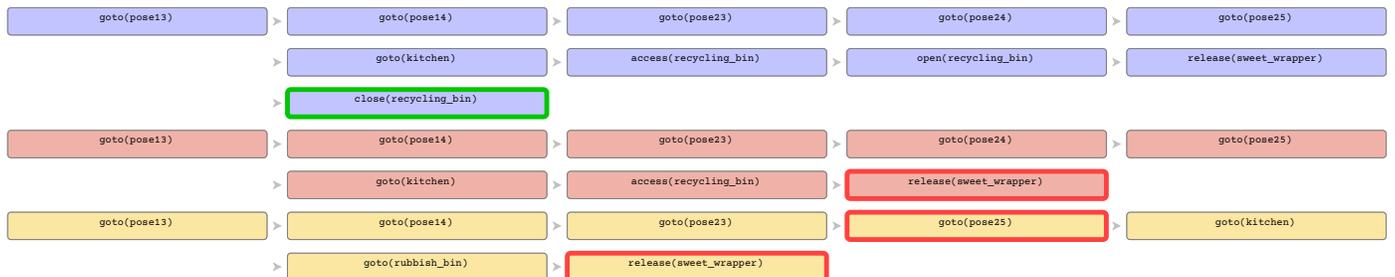
Heat up the chicken kebab.



Something is smelling in the kitchen. Dispose of it.



Throw what the agent is holding in the bin.





Heat up the noodles in the fridge, and place it somewhere where I can enjoy it.



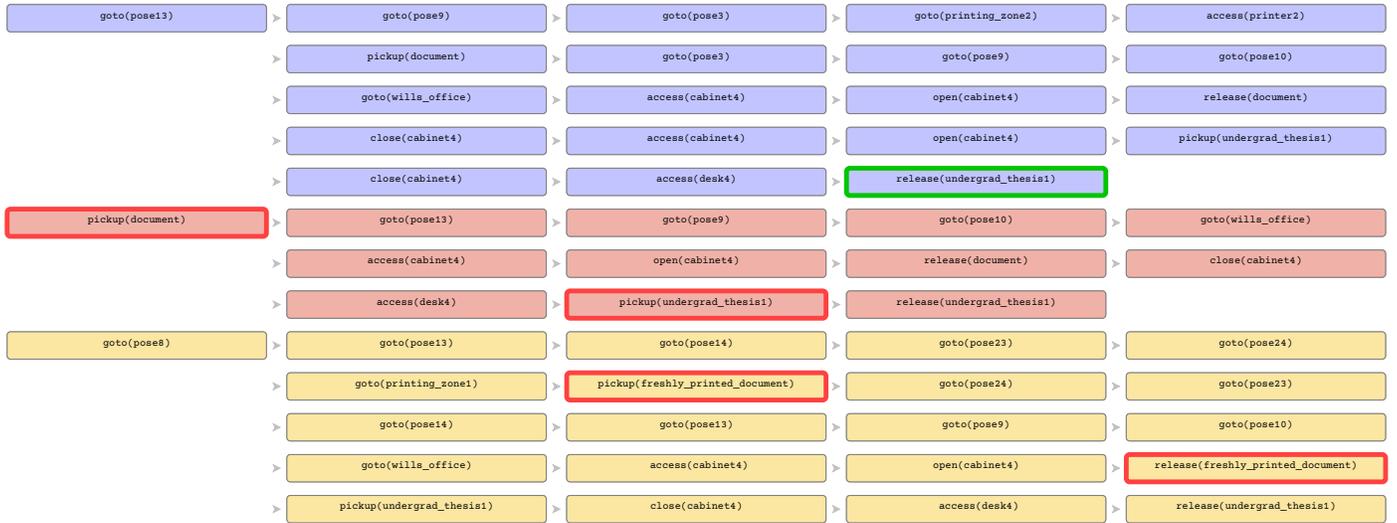
Throw the rotting fruit in Dimity's office in the correct bin.



Wash all the dishes on the lunch table. Once finished, place all the clean cutlery in the drawer.

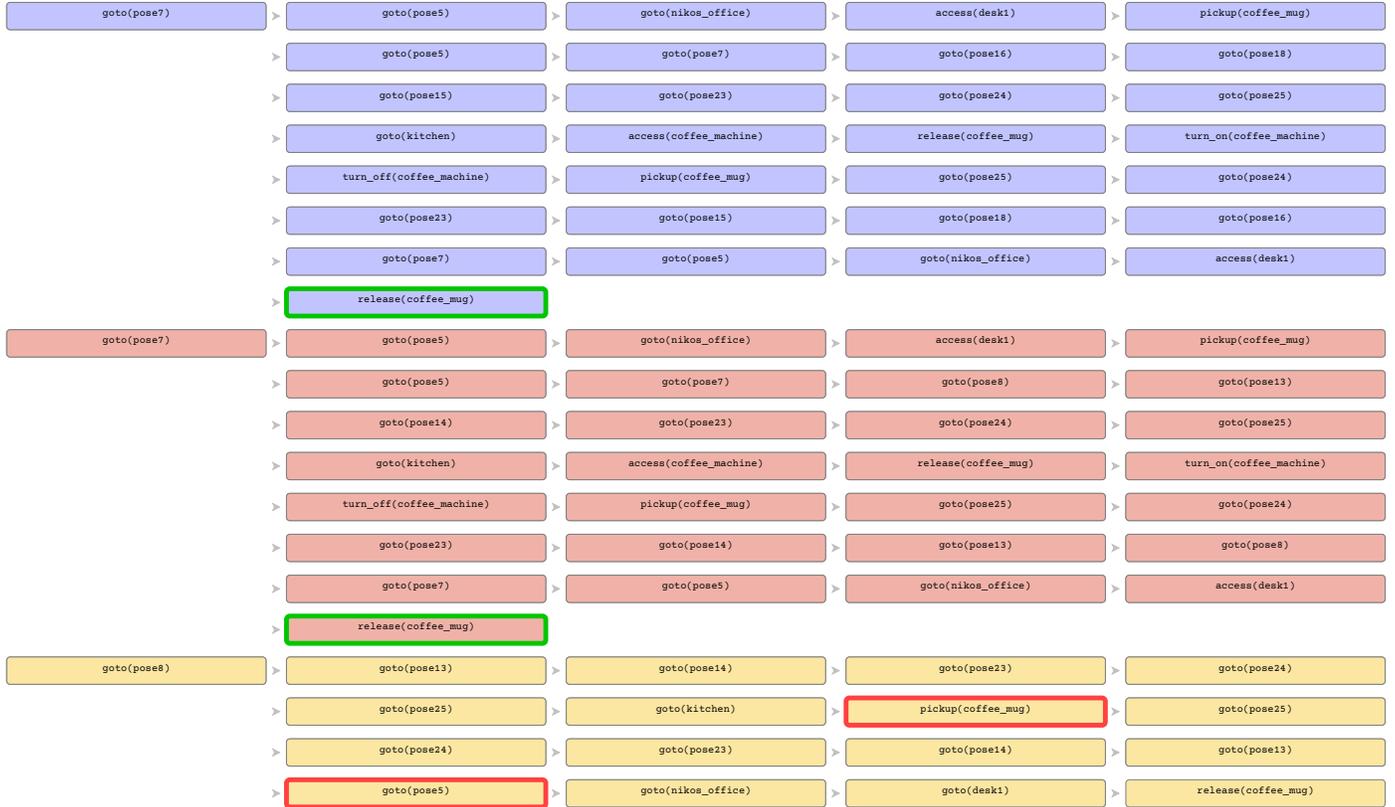


Safely file away the freshly printed document in Will's office then place the undergraduate thesis on his desk.





Make Niko a coffee and place the mug on his desk.

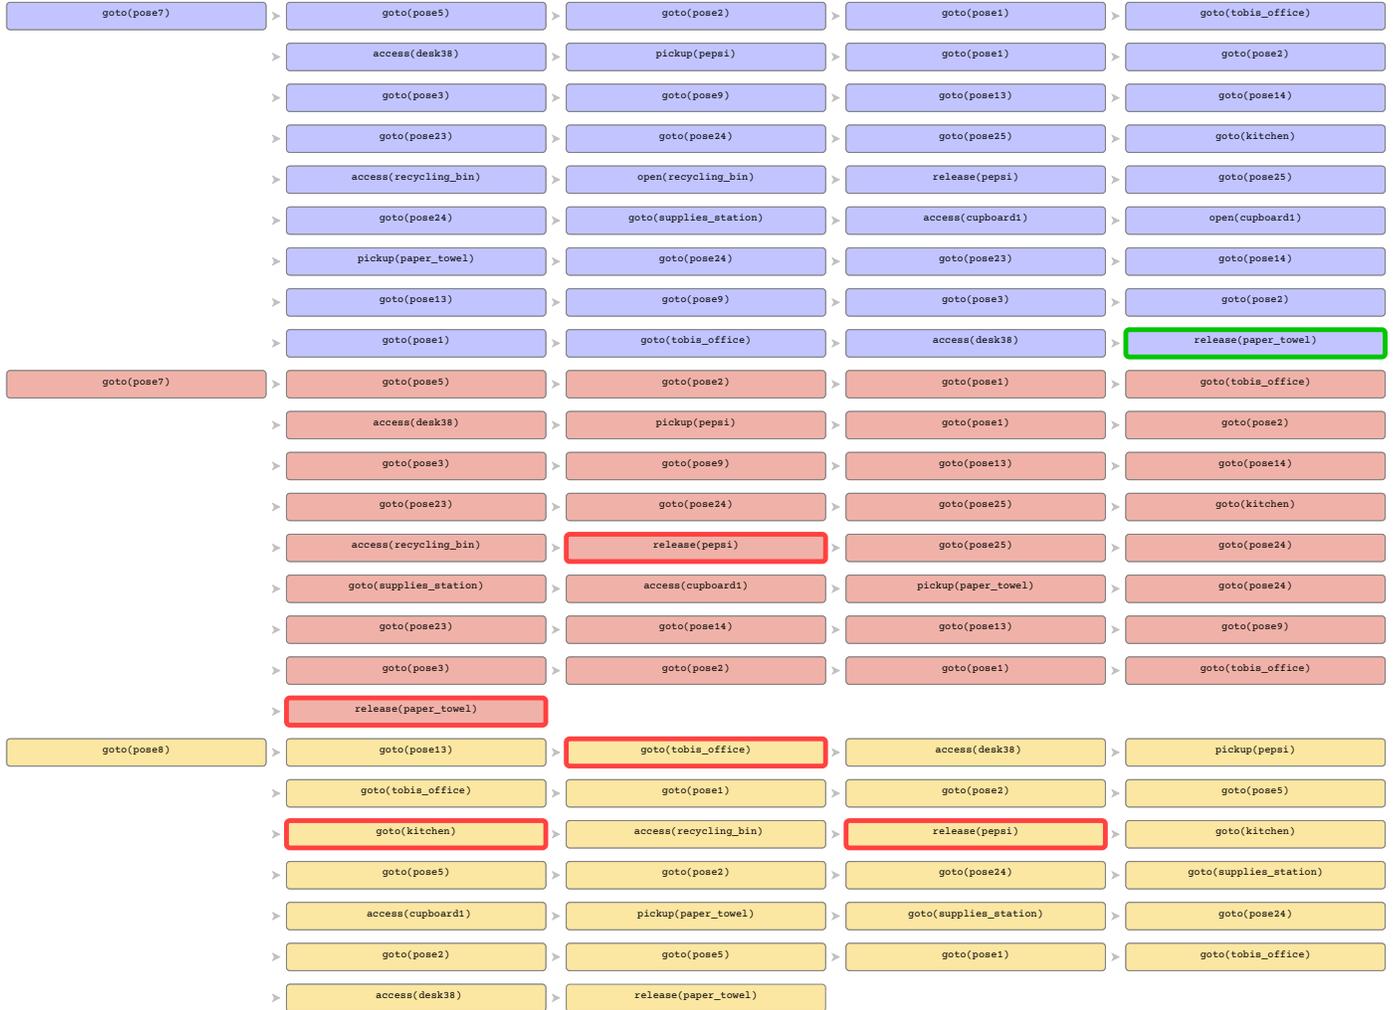


Someone has thrown items in the wrong bins. Correct this.



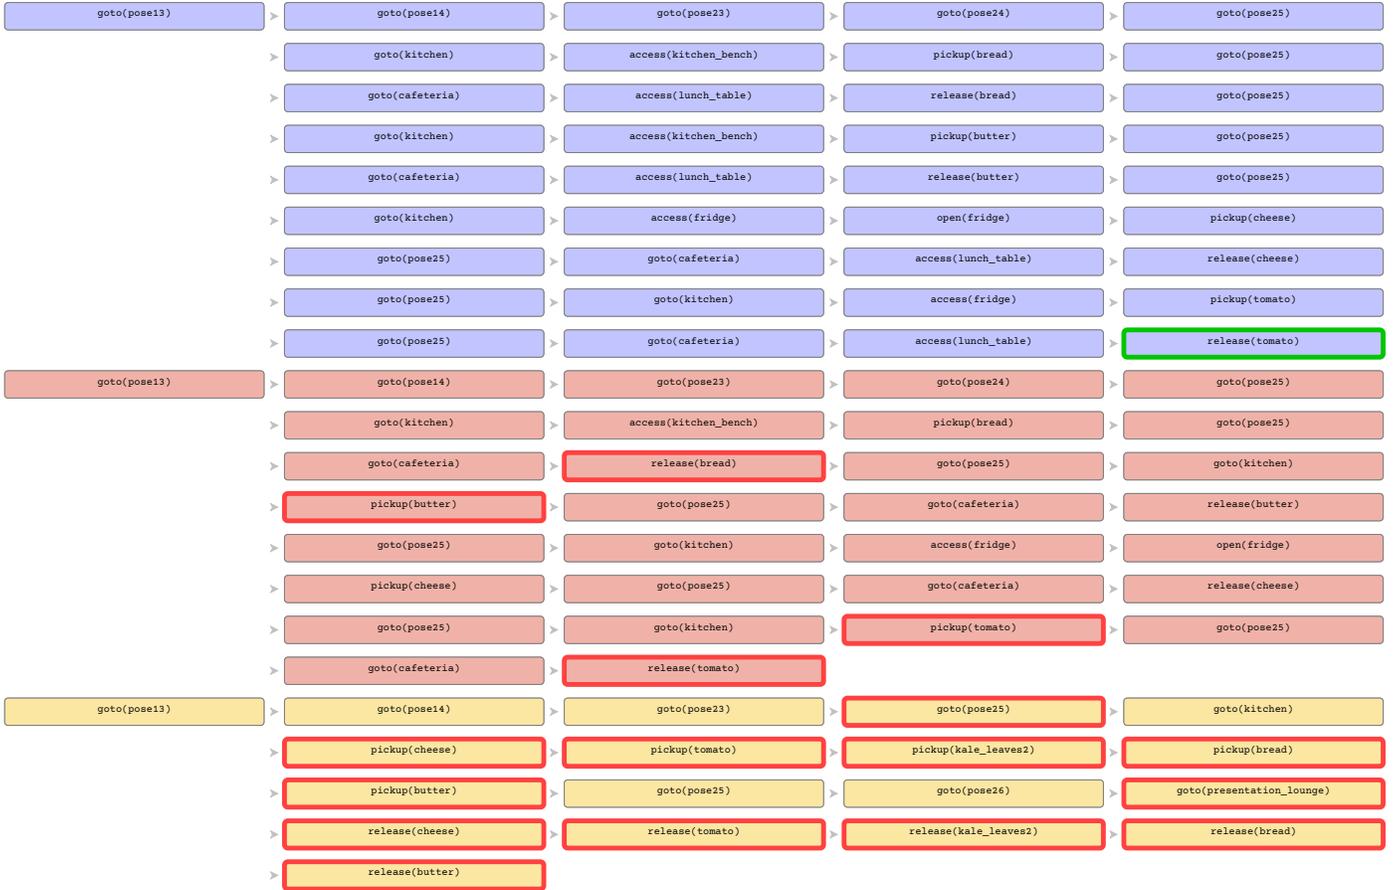


Tobi spilt soda on his desk. Throw away the can and take him something to clean with.

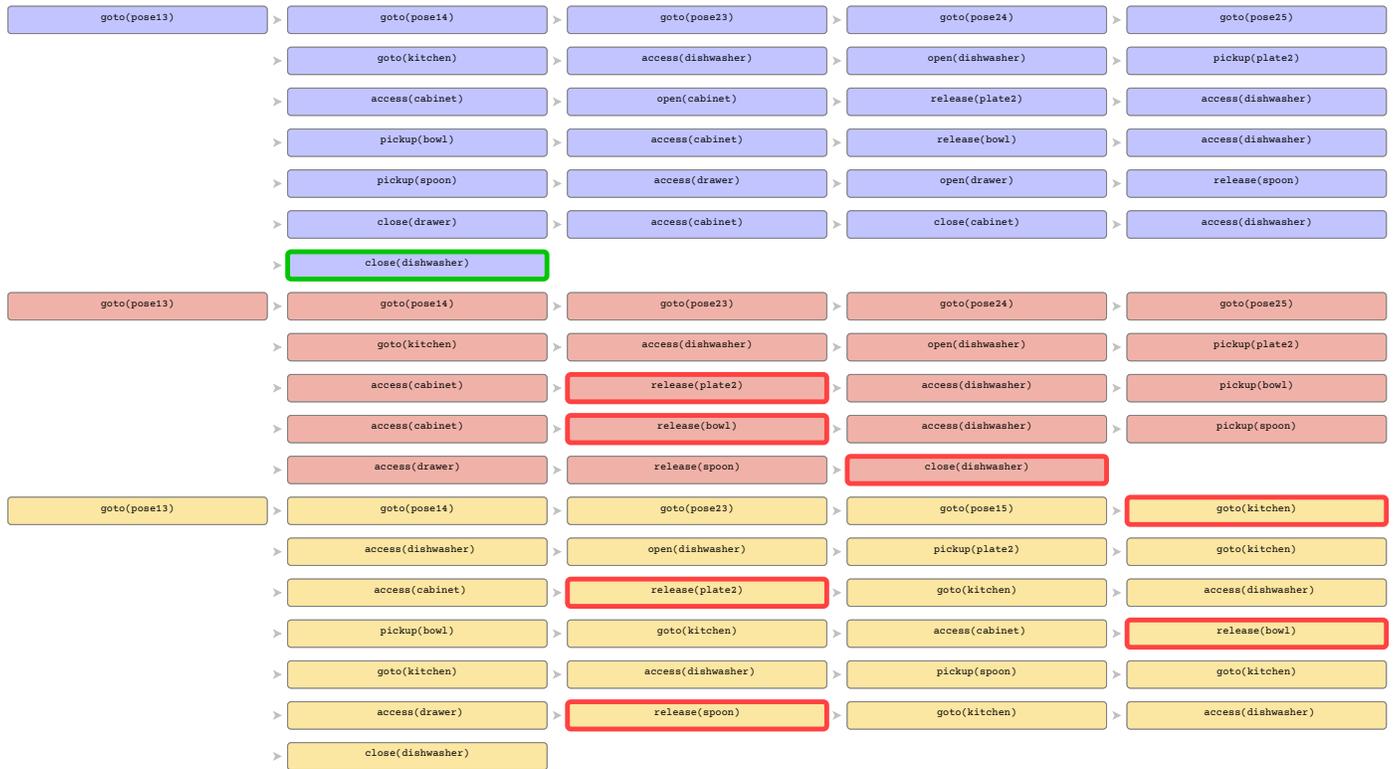




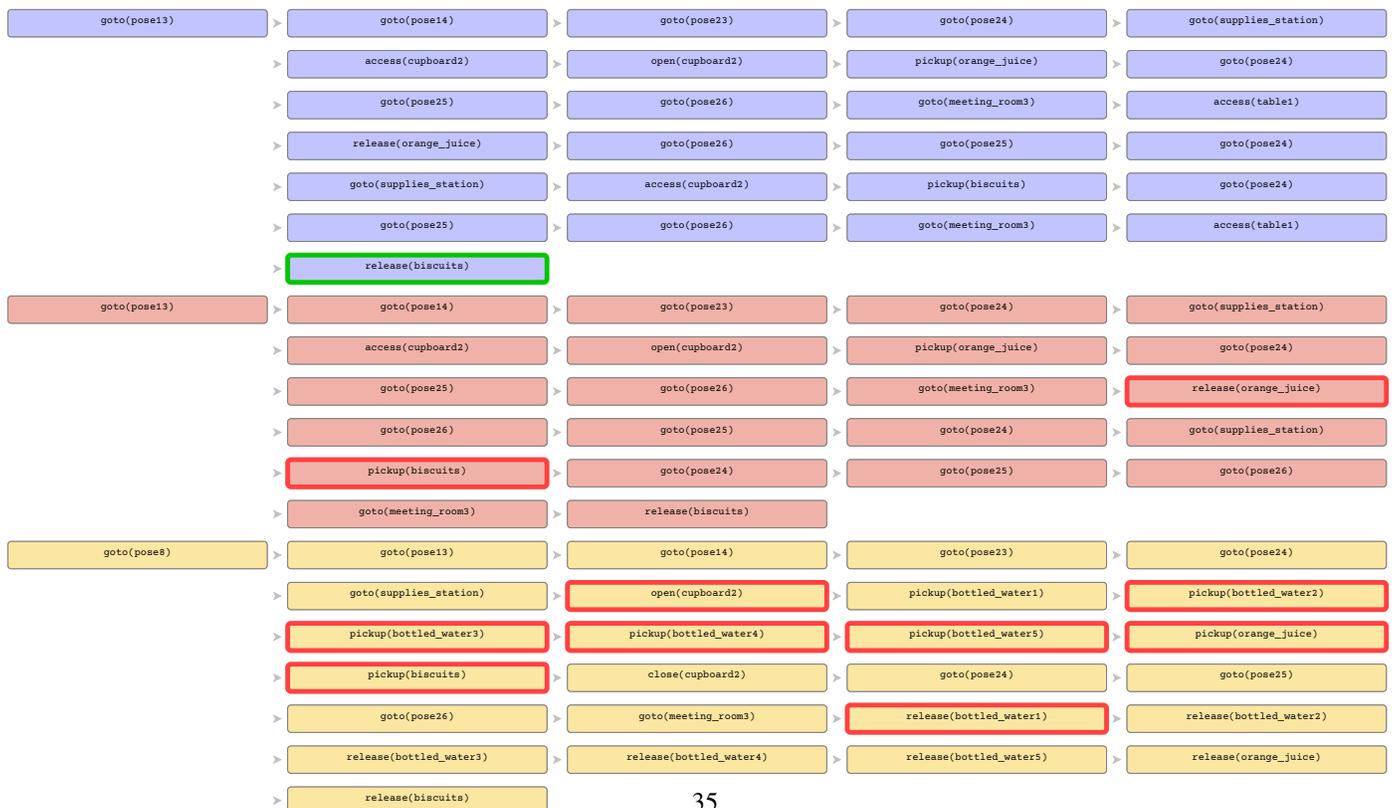
I want to make a sandwich. Place all the ingredients on the lunch table.



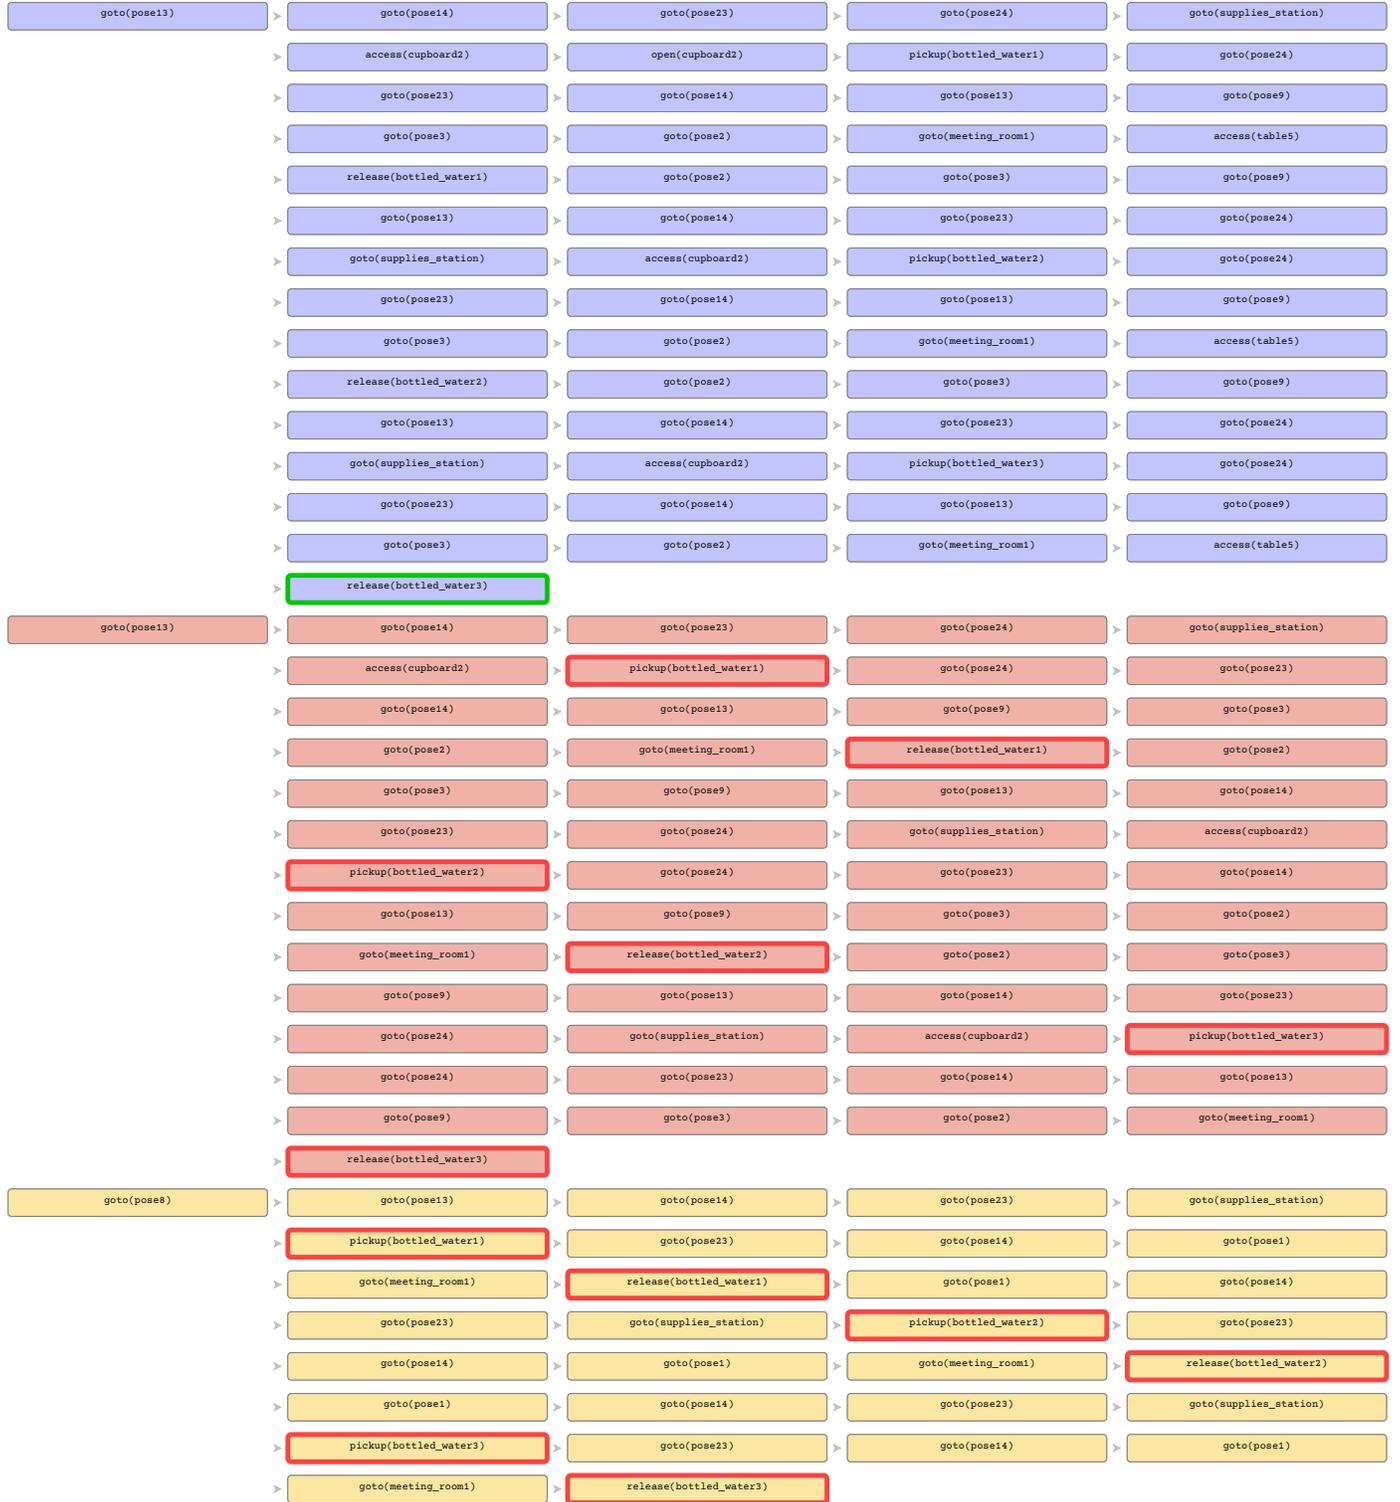
Empty the dishwasher. Place all items in their correct locations



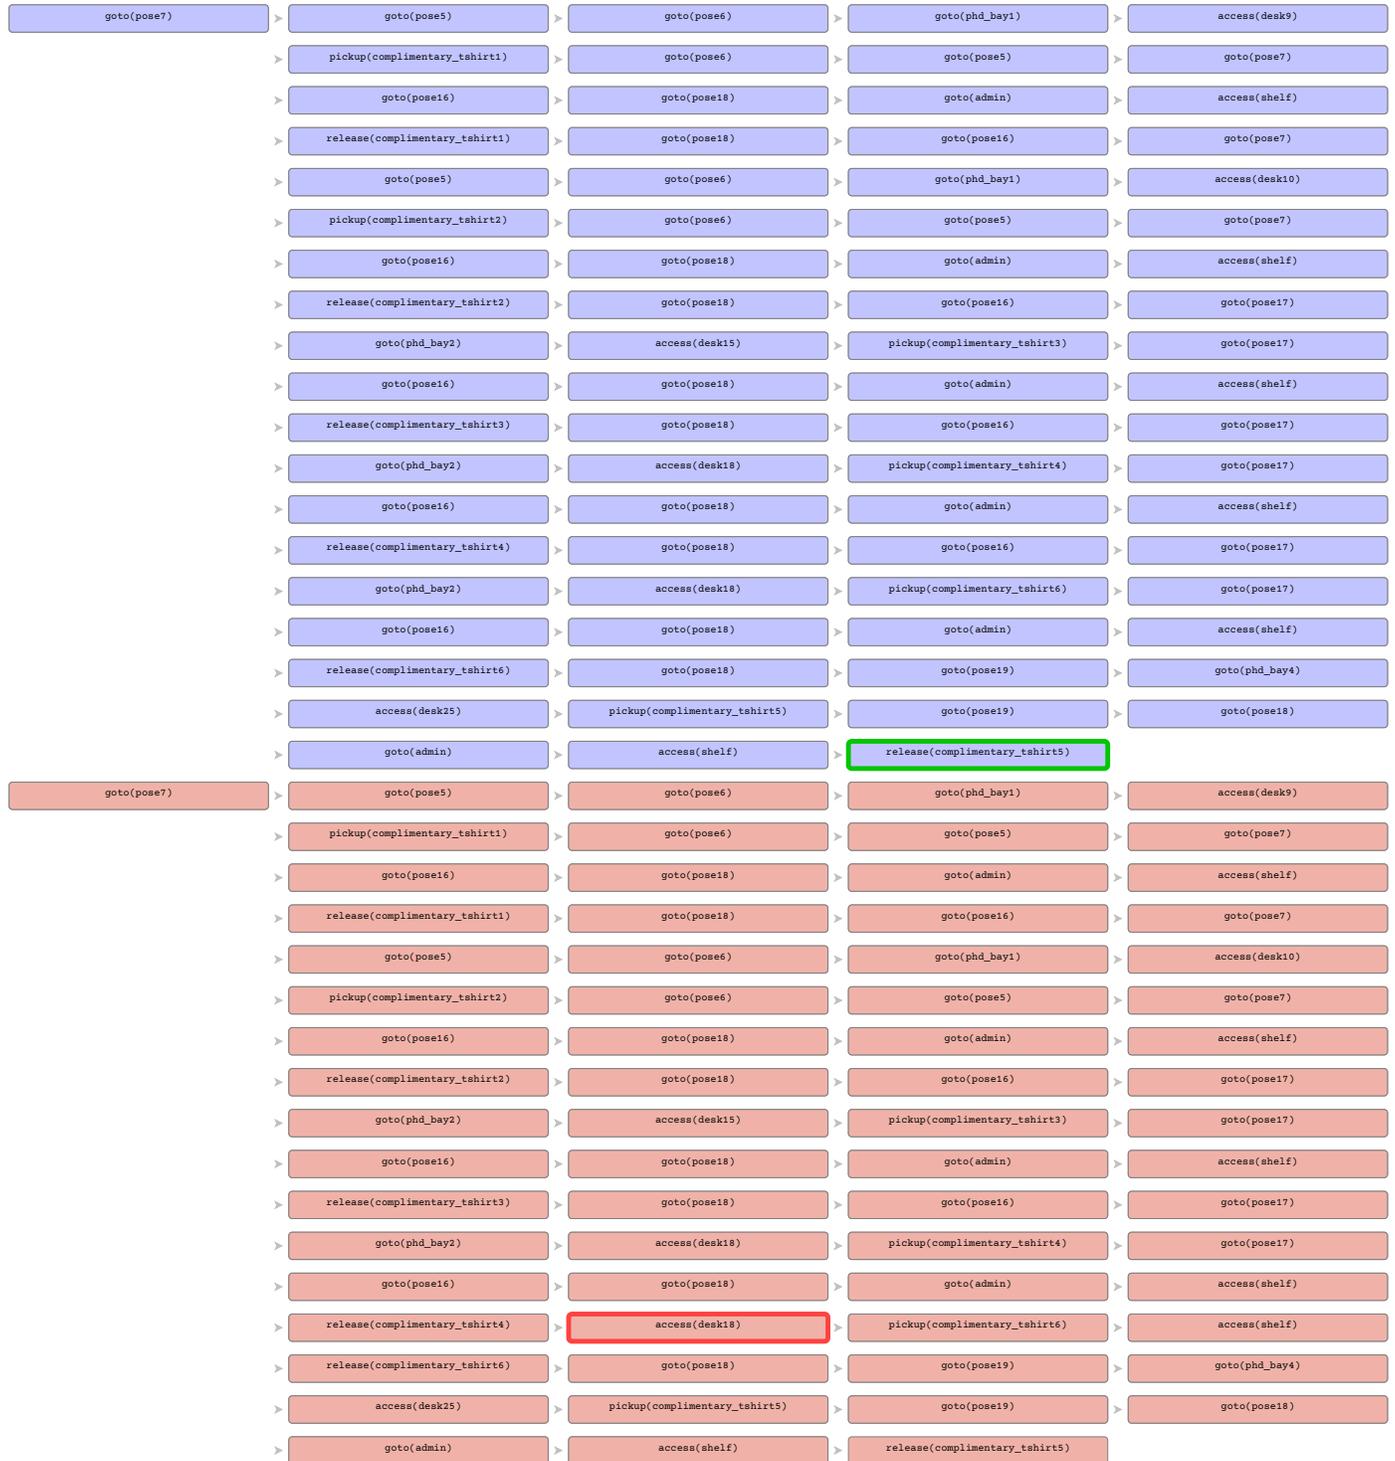
A delegation of project partners is arriving soon. We want to serve them snacks and non-alcoholic drinks. Prepare everything in the largest meeting room. Use items found in the supplies room only.



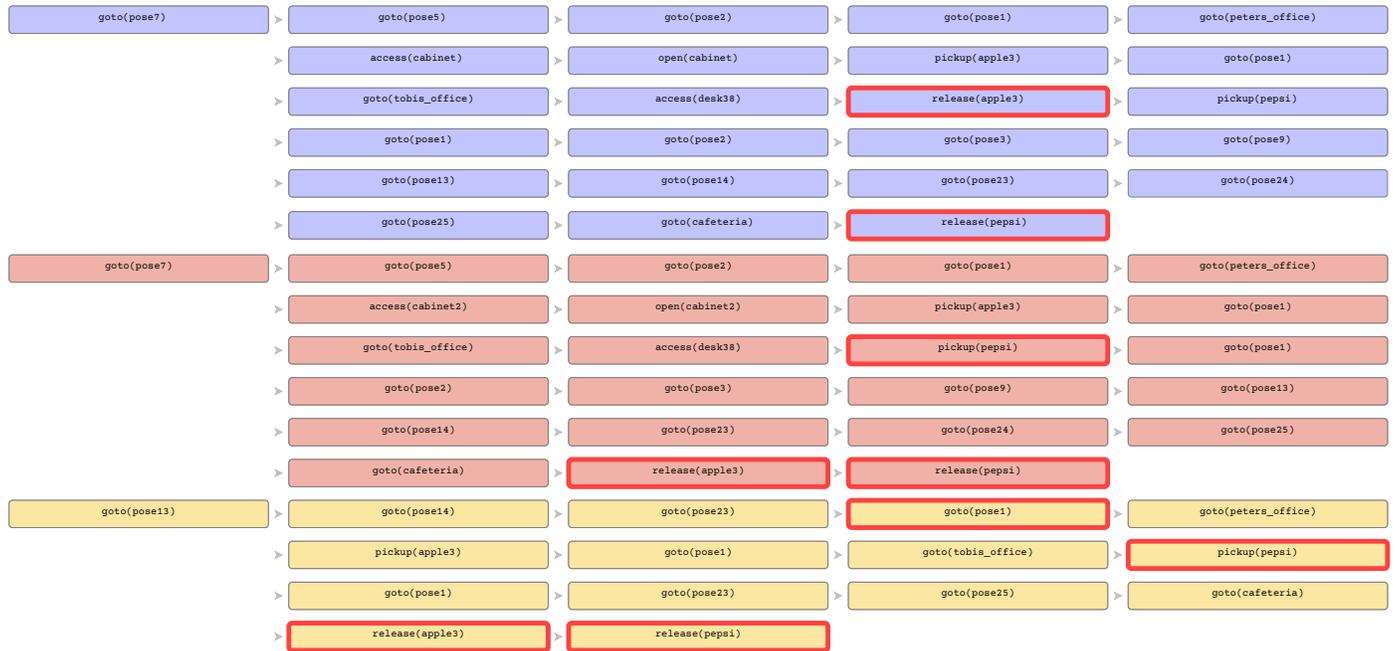
Serve bottled water to the attendees who are seated in meeting room 1. Each attendee can only receive a single bottle of water.



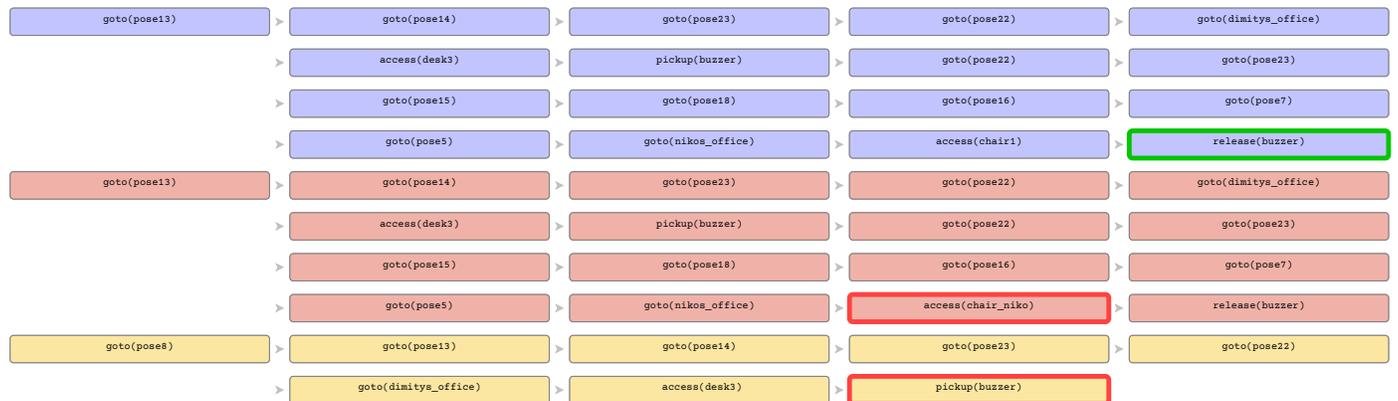
Locate all 6 complimentary t-shirts given to the PhD students and place them on the shelf in admin.



I'm hungry. Bring me an apple from Peter and a pepsi from Tobi. I'm at the lunch table.



Let's play a prank on Niko. Dimity might have something.



514 H Input Prompt Structure

515 Input prompt passed to the LLM for SayPlan. Note that the components highlighted in violet
516 represent static components of the prompt that remain fixed throughout both the semantic search
517 and iterative re-planning phases of SayPlan.

Agent Role: You are an excellent graph planning agent. Given a graph representation of an environment, you can explore the graph by expanding nodes to find the items of interest. You can then use this graph to generate a step-by-step task plan that the agent can follow to solve a given instruction.

Environment Functions:

- goto*(*<pose>*): Move the agent to any room node or pose node.
- access*(*<asset>*): Provide access to the set of affordances associated with an asset node and its connected objects.
- pickup*(*<object>*): Pick up an accessible object from the accessed node.
- release*(*<object>*): Release grasped object at an asset node.
- turn_on/off*(*<object>*): Toggle object at agent's node, if accessible and has affordance.
- open/close*(*<asset>*): Open/close asset at agent's node, affecting object accessibility.
- done*(): Call when the task is completed.

Environment State:

- ontop_of*(*<asset>*): Object is located on *<asset>*
- inside_of*(*<asset>*): Object is located inside *<asset>*
- inside_hand*: Object is currently being grasped by the robot/agent
- closed*: Asset can be opened
- open*: Asset can be closed or kept open
- on*: Asset is currently on
- off*: Asset is currently off
- accessible*: The object is not accessible if it is inside an asset and the asset state is "closed".

Environment API:

- expand_node*(*<node>*): Reveal assets/objects connected to a room/floor node.
- contract_node*(*<node>*): Hide assets/objects, reducing graph size for memory constraints.
- verify_plan*(): Verify generated plan in the scene graph environment.

Output Response Format:

- chain_of_thought*: break your problem down into a series of intermediate reasoning steps to help you determine your next command,
- reasoning*: justify why the next action is important
- mode*: "exploring" OR "planning"
- command*: {"command_name": Environment API call
"node_id": node to perform an operation on
"plan": task plan if in planning mode}}

Example: <see Appendix I and J>

Instruction: Natural language description of the task

3D Scene Graph: Text-serialised JSON description of a 3D scene graph

Memory: History of previously expanded nodes

Feedback: External textual feedback from scene graph simulator

Ensure the response can be parsed by Python `json.loads`.

518 **I Example Input-Output Interaction Between LLM and Scene Graph**
519 **Simulator During Semantic Search**

```
Instruction: make a coffee for Tom and place it in his room
3D Scene Graph: {nodes: {room: [{id: bobs_room}, {id: toms_room}, {id:
jacks_room}, {id: kitchen}, {id: livingroom}], pose: [{id: pose1},
{id: pose2}, {id: pose3}, {id: pose4}, {id: pose5}], agent: [{location:
bobs_room, id: agent}]}, links: [bobs_room↔pose1, bobs_room↔agent,
toms_room↔pose1, toms_room↔pose2, toms_room↔pose5, jacks_room↔pose2,
jacks_room↔pose3, kitchen↔pose3, kitchen↔pose4, kitchen↔pose5,
livingroom↔pose4]}
Memory: []
SayPlan [Mode->Search]: chain-of-thought: explore graph for coffee mug,
coffee machine and somewhere to place the coffee -> perform common sense
exploration first -> explore toms room -> explore kitchen -> if items not
found then explore other nodes
reasoning: i will expand toms room
command: expand(toms_room)
Scene Graph Simulator: <Updated 3D scene graph>
Memory: [toms_room]
SayPlan [Mode->Search]: chain-of-thought: i have found a wardrobe in tom's
room -> leave this node expanded -> the coffee mug is not in his room -> still
have not found the coffee machine -> kitchen might have coffee machine and
coffee mug -> explore this node next
reasoning: i will expand the kitchen next
command: expand(kitchen)
Scene Graph Simulator: <Updated 3D scene graph>
Memory: [toms_room, kitchen]
SayPlan [Mode->Search]: chain-of-thought: i've found the coffee machine in
the kitchen -> the coffee mug however is not in the kitchen -> i will explore
other unexpanded nodes now for the coffee mug -> jacks room, bobs room and the
living room are unexplored yet -> i will try jack's room next
reasoning: i will expand jacks room next
command: expand(jacks_room)
Scene Graph Simulator: <Updated 3D scene graph>
Memory: [toms_room, kitchen, jacks_room]
SayPlan [Mode->Search]: chain-of-thought: jacks room does not have anything
relevant towards solving this task -> I should contract this node to save
token memory -> i will explore other unexplored nodes after doing this
reasoning: contract jack's room to save memory
command: contract(jacks_room)
Scene Graph Simulator: <Updated 3D scene graph>
Memory: [toms_room, kitchen, jacks_room]
SayPlan [Mode->Search]: chain-of-thought: i have found a wardrobe in tom's
room and coffee machine in the kitchen -> the coffee mug is still required for
the task -> explore a new room next
reasoning: i will expand bob's room as it is unexplored
command: expand(bobs_room)
Scene Graph Simulator: <Updated 3D scene graph>
Memory: [toms_room, kitchen, jacks_room, bobs_room]
SayPlan [Mode->Search]: chain-of-thought: i have found the coffee mug, the
coffee machine and tom's wardrobe on the graph -> search complete -> I will
switch to planning mode-> collect coffee mug -> generate plan for making
coffee -> place coffee mug on Tom's wardrobe
reasoning: I will generate a task plan using the identified subgraph
```

520
521

J Example Input-Output Interaction Between LLM and Scene Graph Environment During Iterative Re-planning

```
Instruction: make a coffee for Tom and place it in his room
3D Scene Graph: {nodes: {room: [{id: bobs_room}, {id: toms_room}, {id: jacks_room}, {id: kitchen}, {id: livingroom}], pose: [{id: pose1}, {id: pose2}, {id: pose3}, {id: pose4}, {id: pose5}], agent: [{location: bobs_room, id: agent}], asset: [{room: toms_room, state: free, affordances: [release], id: bed2}, {room: toms_room, state: closed, affordances: [open, close, release], id: wardrobe2}, {room: kitchen, state: closed, affordances: [open, close, release], id: fridge}, {affordances: [turn_on, turn_off], state: off, id: coffee_machine}, {room: bobs_room, state: free, affordances: [release], id: bed1}, {room: bobs_room, state: closed, affordances: [open, close, release], id: wardrobe1}], object: [{affordances: [pickup], state: inside_of(wardrobe1), attributes: "blue", id: coffee_mug}]}, links: [bobs_room↔pose1, bobs_room↔agent, bobs_room↔bed1, bobs_room↔wardrobe1, toms_room↔pose1, toms_room↔pose2, toms_room↔pose5, toms_room↔bed2, toms_room↔wardrobe2, jacks_room↔pose2, jacks_room↔pose3, kitchen↔pose3, kitchen↔pose4, kitchen↔pose5, kitchen↔fridge, kitchen↔coffee_machine, livingroom↔pose4, wardrobe1↔coffee_mug]}
Memory: [toms_room, kitchen, jacks_room, bobs_room]
SayPlan [Mode->Planning]: chain-of-thought: "i have found the coffee mug, the coffee machine and tom's wardrobe on the graph -> search complete -> I will switch to planning mode-> collect coffee mug -> generate plan for making coffee -> place coffee mug on Tom's wardrobe"
reasoning: I will generate a task plan using the identified subgraph
plan: [goto(bobs_room) > access(wardrobe1) > pickup(coffee_mug) > goto(kitchen) > access(coffee_machine) > release(coffee_mug) > turn_on(coffee_machine) > turn_off(coffee_machine) > pickup(coffee_mug) > goto(toms_room) > access(wardrobe2) > release(coffee_mug) > done]
Path Planner:plan: [goto(pose3) > goto(pose4) > goto(pose1) > goto(bobs_room) > access(wardrobe1) > pickup(coffee_mug) > goto(pose2) > goto(pose5) > goto(kitchen) > access(coffee_machine) > release(coffee_mug) > turn_on(coffee_machine) > turn_off(coffee_machine) > pickup(coffee_mug) > goto(pose2) > goto(toms_room) > access(wardrobe2) > release(coffee_mug) > done]
Scene Graph Simulator: "coffee mug is not accessible"
Memory: [toms_room, kitchen, jacks_room, bobs_room]
SayPlan [Mode->Planning]: chain-of-thought: "the coffee mug was not accessible -> the wardrobe it is in is currently closed -> I update the plan to open the wardrobe first -> generate plan for making coffee -> place coffee mug on Tom's wardrobe"
reasoning: I will generate a task plan using the identified subgraph and the feedback from the simulator
plan: [goto(bobs_room) > access(wardrobe1) > open(wardrobe1) > pickup(coffee_mug) > goto(kitchen) > access(coffee_machine) > release(coffee_mug) > turn_on(coffee_machine) > turn_off(coffee_machine) > pickup(coffee_mug) > goto(toms_room) > access(wardrobe2) > release(coffee_mug) > done]
Path Planner:plan: [goto(pose3) > goto(pose4) > goto(pose1) > goto(bobs_room) > access(wardrobe1) > open(wardrobe1) > pickup(coffee_mug) > goto(pose2) > goto(pose5) > goto(kitchen) > access(coffee_machine) > release(coffee_mug) > turn_on(coffee_machine) > turn_off(coffee_machine) > pickup(coffee_mug) > goto(pose5) > goto(pose2) > goto(toms_room) > access(wardrobe2) > release(coffee_mug) > done]
Scene Graph Simulator: Plan Verified
```