


AI for Autonomous Robots: Bridging Theory and Practice

Day 4: Task Planning

Dr. Timothy Wiley


School of Computing Technologies
RMIT University



—
ESSAI July 2026

Motivation

—
ESSAI July 2026



Find a Kitchen Item

Consider the following problem for a service robot:

- *“Human says to robot: Go a fetch a can of <beverage> from the fridge and bring it to me in the lounge room”*

Questions:

- What type of information do we need to know for this problem?
- What type of information should be represented at the Deliberative layer?
- What type of "deliberations" need to be made for planning?





--- An Observation

The summer school have a variety of courses on:

- Knowledge Representation and Reasoning
- Symbolic Logics
- Symbolic Planning

What will be considered in this course is the *practical* issues of deploying a symbolic planner onto a robotic platform




Software Architectures

Recap

—
ESSAI July 2026




Software Architectures: Three-Layer



Bonasso, P. et. al. (1997). Experiences with an architecture for intelligent, reactive agents. Journal of Experimental & Theoretical Artificial Intelligence 9(2-3):237– 256



Software Architectures: RCS



Albus, J. S. & Barbera, A. J. (2005) RCS: A cognitive architecture for intelligent multi-agent systems. Annual Rev Control 29, 87–99.



Symbolic Task Planning

For robotics


—
ESSAI July 2026



Task Planning formulation

The discussions here define symbolic task planning as a symbolic state-action plan, with:


- A State, s , that are symbolic representations of the numeric state space of the robot, *and*, the environment
- An Action, a , are symbolic representations of robot movement.
 - Note that actions do not necessarily need to directly correlate to one actuator movement
- Actions are temporal, and may take a variable length of time to execute



Task Planning formulation

Therefore a *task plan* is:


- A sequence of actions that, when executed, enable the robot so change from an initial state to a goals state.
- Actions connect intermediate states



Executing Task Plans

A task plan is (naively) executed by:

1. Executing the first action of the task plan until the first intermediate state is reached
2. Execution continues with subsequent actions, until the relevant intermediate states are reached
3. Execution terminates on reaching the goal state.



STRIPS

—
ESSAI July 2026





STRIPS Planning (as done on Shakey!)

STRIPS Planning is a very simple method for symbolic planning.

- State: Represented by logic predicates
- Actions:
 - Name
 - Precondition
 - Postcondition
 - Add list
 - Delete list




A note for ESSAI'26

It's likely many are quite familiar with STRIPS planning, and far more complex planning languages.

We are going to start with STRIPS to investigate the issues of the practical side of task planning on robot systems, and motivate more modern planning languages.



Blocks World



States:


- On(A,C)
- Clear(B), Clear(C)
- OnTable(A), OnTable(B)
- Holding(/)

Actions:

- Pickup(X)
- Putdown(X)
- Stack (X,Y)
- Unstack (X,Y)



Blocks World



Pickup(x):


- Preconditions:
 - Clear(x)
 - OnTable(x)
 - Holding($\emptyset$)
- Delete list:
 - Clear(x)
 - OnTable(x)
 - Holding($\emptyset$)
- Add list:
 - Holding(x)

Putdown(x)

- Precondition:
 - Holding(x)
- Delete list:
 - Holding(x)
- Add list:
 - Clear(x)
 - OnTable(x)
 - Holding($\emptyset$)



Blocks World



Unstack(x,y):

- Preconditions:
 - Clear(x)
 - On(x,y)
 - Holding($\emptyset$)
- Delete list:
 - Clear(x)
 - On(x,y)
 - Holding($\emptyset$)
- Add list:
 - Holding(x)
 - Clear(y)

Stack(x,y)

- Precondition:
 - Holding(x)
 - Clear(y)
- Delete list:
 - Holding(x)
 - Clear(y)
- Add list:
 - Clear(x)
 - On(x,y)
 - Holding($\emptyset$)



Baxter w/ Blocks World (Not quite STRIPS)



B. Hengst, et. al, A framework for integrating symbolic and sub-symbolic representations. 25th International Joint Conference on Artificial Intelligence IJCAI-16. New York, New York, USA, 2016.





Issues that STRIPS highlights

STRIPS Planning is useful for highlighting challenges for symbolic planning in real-world robotics:

- “Perfect” state assumption
- Symbolic to sub-symbolic correlation
- Observability / Occlusion
- Durative Actions
- Actions failure
- Outside Influences
- Response to real-time sensor updates
- Parallel Action




Baxter w/ Blocks World (Not quite STRIPS)



Baxter in blocks-world



“Mind’s Eye” physics simulator




A block and end-effector

B. Hengst, et. al, A framework for integrating symbolic and sub-symbolic representations. 25th International Joint Conference on Artificial Intelligence IJCAI-16. New York, New York, USA, 2016.



Symbolic to Sub-Symbolic

—
ESSAI July 2026





State Representation at TLA layers

In a deliberative layer a symbolic representation, such as in blocks-world, may represent the world state by:

- $\text{OnTable}(A)$
- $\text{On}(B,A)$

Below the reactive layer is sub-symbolic representations, that "does not have symbolic entities or discrete elements" but, typically, involves some form of numeric information.

In the reactive layer, this sub-symbolic information are sensor information, including the camera. For the block's world this would be a camera image of the layout of the world





State Representation at TLA layers

An current open research challenge is how to "bridge" these representations in the sequencing layer. This is because, the symbolic world is "idealistic", but the "real-world" has:

- Noise
- Error
- Non-determinism
- Hidden information
- Dynamic changing information





Potential Solutions

There are many potential solutions:

- Improve the "richness" of symbolic representations - that is, the simple modelling is insufficient
- Embed sub-symbolic information in symbolic representations, such as:
 - `OnTable (A, x, y, z)`
 - `On(C,A, <relative position)`


Using symbolic modelling of continuous variables, rather than pure symbolic terms



Tree / Graph

Behaviour Structures

—
ESSAI July 2026





Decision Trees

Decision tree's can be used to define a 'planning' or behaviour generation where the plan of actions to take is determined by the structure of nodes the decision tree.

- Must conform to a tree-structure - that is, no loops
- Operate on a 'tick' cycle, where the tree is re-evaluated entirely from top-to-bottom every tick
- The tree is evaluated until a leaf node is reached
- Nodes may: decide which children node is called, pass data to children nodes, or generate a behaviour to execute in sequencing/reactive layer(s).
- Nodes may: permit parallel actions, contain sub-decision trees, or other planning approaches.



Decision Trees




Image: Röfer, T. RoboCup 2017: Robot World Cup XXI. Lecture Notes Computer Science 135–142 (2018) doi:10.1007/978-3-030-00308-1_11.



Decision Trees




Image: RedbackBots internal documentation 2023.





Decision Trees

Advantages of decision trees:

- Relatively simple structure to understand and implement
- In principle, should be very quick to evaluate
- Good for 'simple' rule-based or deterministic behaviours
- Can avoid being 'stuck' in local decisions, as they 'globally' re-evaluate the decision
- Markovian decision making: can be quick to switch context(s)/decisions with new information





Decision Trees

Disadvantages:

- 'Complex' behaviours can be quite challenging to design
- Grow 'exponentially' in complexity as more nodes as added, in terms of the number of interactions of nodes, and decision between nodes.
- Have no real concept of 'planning'. A plan is 'implicitly represented' in the decision tree structure.
 - Do not 'look ahead' for future decisions in future ticks.
 - No past introspection of previous decisions.
 - No concept of "durative" actions operating over time.





Finite State Machines


Finite State Machines allow for a more explicit concept of a behaviour (or action) operating over the time, until a decision is made to change to a new behaviour (or action).

Finite State Machines also operate on a tick-basis, however compared to decision trees, each tick:

- Determines if the current node should be run, or if there should be a transition to another node.
- If there is no transition, the node:
 - Generates a behaviour to execute in sequencing/reactive layer(s), or
 - Continues executing the previous generated behaviour.



Finite State Machines





Finite State Machines

Advantages:

- Only the current node of the FSM is evaluated.
- Quick to evaluate and execute.
- Prefer smooth transitions between adjacent nodes, rather than abrupt changes to completely different behaviours.
- Concept of 'durative' actions.
- Markovian Node transitions:
 - Once a transition is made, there is no history at why that node is now 'running'.
 - No complex interaction between 'chains' of node transitions.
- Prefer simple structure with few transitions between nodes.





Finite State Machines

Disadvantages:

- Logic might need to be repeated in different nodes where there are similarities.
- The lack of a global view can lead to becoming 'stuck' in local regions of the state machine or be slower 'to quickly react' to major changes in the environment.
- Also have no real concept of 'planning'.
- Difficult to define complex behaviours that require transitions across sequences of nodes.
- The graph may devolve into a complex structure, at worse to a complete (fully connected) graph.





Behaviour Tree

Behaviour trees are a mix of the tree structure of a decision tree, and the durative operation of finite state machines.

Behaviour tree are a tree like-structure of Control Flow Nodes, and Execution Nodes:

- Control nodes determine how execution nodes are structured.
- Execution nodes generate behaviours for the sequential/reactive layer.





Behaviour Tree

An execution node:

- Operates durativity each tick until it either terminates in success or failure.
- Returns each tick one of three statuses to its parent node:
 - Running - the node is executing
 - Success - the node has terminated in success
 - Failure - the node has terminated in failure

If the root node of the behaviour tree returns failure, the whole tree may be re-executed from the start. Additionally, the whole tree may be intermittently re-evaluated/re-started to avoid becoming stuck in a local execution node.





Behaviour Tree


Behaviour trees provide two core control nodes:

- Sequence
 - A sequential order for nodes to operate.
 - Presumes each node succeeds.
- Fallback
 - A sequence of alternative options of nodes to execute if a node fails

Behaviours Trees are used to implement the ROS2 Navigation (Nav2) stack.



Behaviour Tree





Behaviour Tree

Advantages:

- Only the current node is evaluated.
- Quick to evaluate and execute.
- Concept of 'durative' action.
- Provides a more "programming" control flow structure over node execution.
- The structure can be thought of in theory as a pre-defined plan.
- Allows for 'planning for failure'.





Behaviour Tree

Disadvantages:

- Can become 'stuck' in execution nodes.
- May not react 'quickly' to major changes in the environment.
- Also have no real concept of 'planning'.
- Do not handle re-planning very well, outside of what is explicitly represented in the behaviour tree for failure handling.
- Unable to return to earlier items in the tree sequence without the whole tree first failing. This can lead to unintended behaviour if later execution nodes commence operation after changes in the environment break assumptions of earlier nodes.



Agent Programming Languages

GOLOG

—
ESSAI July 2026



GOLOG Programming Language

Golog is a logic-based agent programming language built on-top of Prolog, that facilitates the representation and execution of actions:

1. **Actions:** Predicates that can be any operation that modifies the state of the environment or the agent itself.
2. **Fluents:** represent state information that can change over time. This allows reasoning about dynamic aspects of the environment and update the state information.
3. **Concurrent Actions:** execution of multiple actions concurrently.
4. **Non-Deterministic Choice:** This allows the agent to determine different possible actions where it is unknown the probability of either alternative to explore different possibilities.

Goals: represent the desired state of the agent





Task Planning Problem

Devise a task plan to *“go and get a beverage from the fridge”*.

Assumes knowledge of:

- Rooms of the house
- Objects in each room
- Relationship between objects
- Tracking of “person” locations
- Tracking state of “modifiable” objects



GOLOG Program: Fluents

```
/* Robot Navigation in a House using Golog
*
* This program represents a simple domain where a robot can move
* between different rooms in a house: kitchen, lounge, bedroom, and hallway.
*/

/* Fluents (properties that can change with actions) */

% The robot's location is a fluent
prim_fluent(robotLoc(Loc)).
```



GOLOG Program: Actions

```
/* Actions */
```

```
% move(FromLoc, ToLoc): The robot moves from FromLoc to ToLoc  
prim_action(move(FromLoc, ToLoc)).
```

```
/* Action Preconditions */
```

```
% To move from one location to another, the robot must be at the first location  
% and the locations must be connected  
poss(move(FromLoc, ToLoc), and(robotLoc(FromLoc), connected(FromLoc, ToLoc))).
```

```
/* Successor State Axioms */
```

```
% How the robot's location changes  
causes_val(move(_FromLoc, ToLoc), robotLoc(_), ToLoc).
```



GOLOG Program: Domain Knowledge

```
/* Domain Knowledge */  
  
% Connected rooms (bi-directional connections)  
connected(hallway, kitchen).  
connected(kitchen, hallway).  
  
connected(hallway, lounge).  
connected(lounge, hallway).  
  
connected(hallway, bedroom).  
connected(bedroom, hallway).
```



GOLOG Program: Procedures

```
/* Procedures */
```

```
% visit(Loc): visit a specific location
proc(visit(Loc),
[?(robotLoc(CurrentLoc)),
if(CurrentLoc = Loc,
?(true), % Already at the location, do nothing
findPath(CurrentLoc, Loc, Path) : followPath(Path))
]).
```

```
% followPath(Path): follow a path of locations
proc(followPath([]), ?(true)). % Empty path, do nothing
proc(followPath([Next|Rest]),
[?(robotLoc(Current)),
move(Current, Next),
followPath(Rest)]).
```



GOLOG Program: Procedures

```
/* Path Planning */
```

```
% findPath(Start, Goal, Path): find a path from Start to Goal  
% This is a simplified version that assumes the house topology  
findPath(Start, Start, []). % Already at the destination  
findPath(Start, Goal, [Goal]) :- connected(Start, Goal).  
findPath(Start, Goal, [Next|Rest]) :-  
    connected(Start, Next),  
    Next \= Goal,  
    findPath(Next, Goal, Rest).
```



GOLOG Program: Initial State & Goal

```
/* Initial State */
```

```
% The robot starts in the hallway  
initially(robotLoc(hallway)).
```

```
/* Example Queries */
```

```
% Example query to visit the kitchen
```

```
do(visit(kitchen), S0, FinalState).
```

```
% Find a plan to achieve a goal  
plan(robotLoc(kitchen), Plan).
```



GOLOG Program: Main planning program

```
main :-  
    goal,      % Define the goal  
    choose_fridge(Fridge),  
    fridge(Location, Fridge),  
    goto(Location),  
    open_and_grab(coke).  
  
% Initial State  
at(starting_location).  
connected(starting_location, kitchen).  
connected(kitchen, garage).  
fridge(kitchen, fridge1).  
fridge(garage, fridge2).  
item_in_fridge(fridge1, coke).
```



LLMs & Robot Behaviour Execution

—
ESSAI July 2026





— Naively

Naively we can generate robot (ROS2) code with LLMs.

Worked example with Claude. AI


What do you want to get the robot to do?



How or Why to use an LLM?

Some items to ponder:

- At what layer of a Robot Software Architecture should an LLM be leveraged?
- LLMs are (currently) poor at planning and sequential/temporal reasoning.
- What sensing and world modelling data/inputs can be leveraged?
 - Text-only?
 - Multi-model?
- What are an LLMs strengths?
- How efficient is an LLMs use?




LLM for Behaviour Generation

The most obvious, and currently typical, use of LLMs is for behaviour generations.

But at what level?


- Reactive control commands are too slow
- Deliberative planning is perhaps possible (if good plans are known), but current LLMs are poor at planning
- Sequencing level skill acquisition, translating a 'high-level goal' into reactive operations is the current place.



LLM for Behaviour Generation

Thus, at the sequencing layer the output of LLM behaviour generation is essentially 'robot code'

- Presume a software API with the reactive layer control
- Generate 'code' to sequence/execute the robot's actions
- The code can bridge both symbolic and sub-symbolic (numeric) representations.




Multi-Modal Input

The sensing (or world model) input to the LLM needs to be considered along with the query of which the LLM will answer.

The most similar sensing on which models have been extensively trained are images

So our query is:

- Given image inputs
- Given a 'task'
- Generate 'code' or a form of an action



RT-1/2: Vision-Language-Action Models

Internet-Scale VQA + Robot Action Data



Q: What is happening in the image?

A: 311 423 170 55 244

A grey donkey walks down the street.

Q: Que puis-je faire avec ces objets?

A: 3455 1144 189 25673

Faire cuire un gâteau.



Q: What should the robot do to <task>?

A: 132 114 128 5 25 156

$\Delta T = [0.1, -0.2, 0]$

$\Delta R = [10^\circ, 25^\circ, -7^\circ]$

Vision-Language-Action Models for Robot Control

Q: What should the robot do to <task>? A: ...



RT-2

Large Language Model

ViT

A: 132 114 128 5 25 156

De-Tokenize


$\Delta T = [0.1, -0.2, 0]$
 $\Delta R = [10^\circ, 25^\circ, -7^\circ]$

Robot Action

Co-Fine-Tune

Deploy


Closed-Loop Robot Control



Put the strawberry into the correct bowl




Pick the nearly falling bag




Pick object that is different



RT-1/2: Vision-Language-Action Models



RT-1/2: Vision-Language-Action Models



Google DeepMind: <https://robotics-transformer2.github.io>



Thank you

Dr. Timothy Wiley
School of Computing Technologies
RMIT University



—
ESSAI July 2026

— Thank you

I would like to express my deep thanks for everyone who has attended these classes over the week. I have very much enjoyed this week. I hope that I have provided some insight into what robotics can offer and how to apply AI in real-world settings.

Please feel free to connect and stay in touch!

- Course Resources: <https://timothy-wiley.github.io/essai.html>
- LinkedIn: <https://www.linkedin.com/in/timothy-wiley-948a9113/>
- Email: timothy.wiley@rmit.edu.au



Noon Gudgin

Thank you

ESSAI July 2026

