Resource-rational Task Decomposition to Minimize Planning Costs

Carlos G. Correa*, **Mark K. Ho***, **Fred Callaway, Thomas L. Griffiths**

*Contributed equally

Princeton University

We propose a normative model of hierarchical task decomposition that chooses subgoals that minimize the computational cost of searching for a plan.

Background

- Search is intractable due to the *curse of dimensionality*[1]. Good problem representations are essential for efficient search[2].
- Task decomposition into hierarchy can yield problem representations to make the search process more efficient[3].
- Previous models of human task decomposition pose task decomposition as inference of latent structure of optimal behavior[4] or the environment[5].

Model

We formalize task decomposition as a *resource-rational representation problem*: What hierarchical representations best support computationally efficient search?

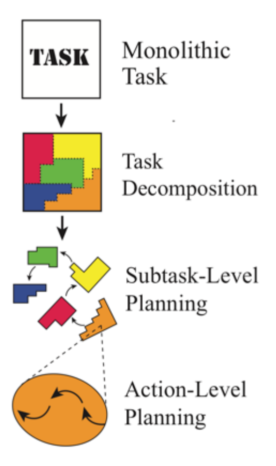

Our model specifies **three nested levels of optimization**:

- **Action-level planning:** find a sequence of actions to reach a state.
- **Subtask-level planning:** select a sequence of subgoals to reach the goal based on the reward and planning cost to get to each subgoal.
- **Task decomposition:** identify a set of subgoals that maximizes subtask-level planning performance.

Action-Level Planner

Breadth-First Search (BFS): Explores states in the order of their distance from the starting state.

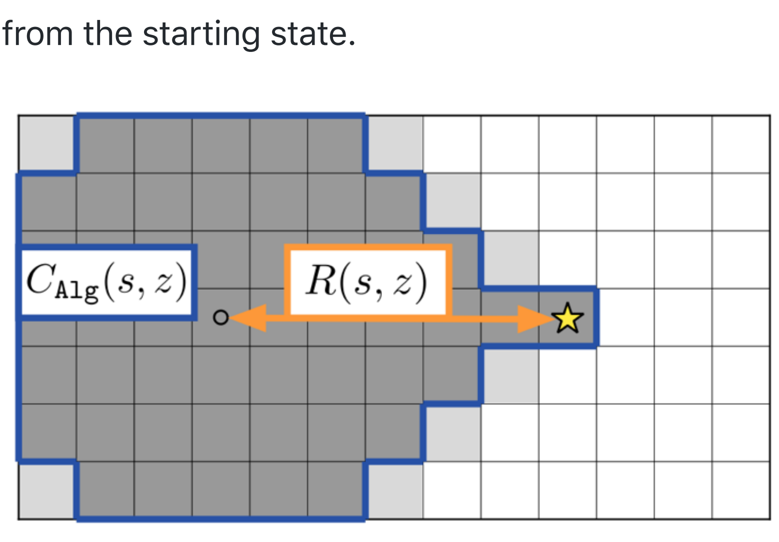

$R(s, z)$ is reward between state s and z. For shortest path problems, this is the negative distance.

$C_{\texttt{Alg}}(s, z)$ is cost to plan from state s to z using algorithm **`Alg`**. For BFS, this is number of states visited by search, which is the number of states closer to s than z is to s.

Subtask-Level Planner

Given a task decomposition $\mathcal{Z}$, we use value iteration to identify the optimal sequence of subtasks $z \in \mathcal{Z}$ to reach a goal g from state s.

$$V_{\mathcal{Z}}^{g}(s) = \max_{z \in \mathcal{Z}} \left\{ R(s, z) - C_{\texttt{Alg}}(s, z) + V_{\mathcal{Z}}^{g}(z) \right\}$$

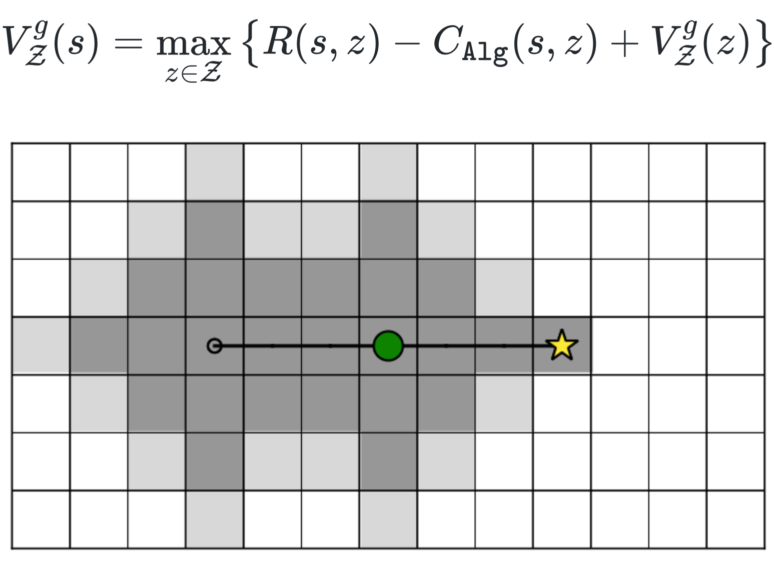

Task Decomposition

We find the set of subgoals $\mathcal{Z}$ that maximizes performance $V_{\mathcal{Z}}^{g}(s)$ given a task distribution over start state s and goal g.

$$\mathcal{Z}^* = \arg\max_{\mathcal{Z}} \mathbb{E}_{s,g}[V_{\mathcal{Z}}^{g}(s)]$$

Results

Tasks are decomposed at bottleneck states, matching intuitions in grid worlds and previous quantitative predictions in other tasks[4].

Grid worlds

Subgoals (left) and subgoal-level policy with ⭐ as goal g (right).

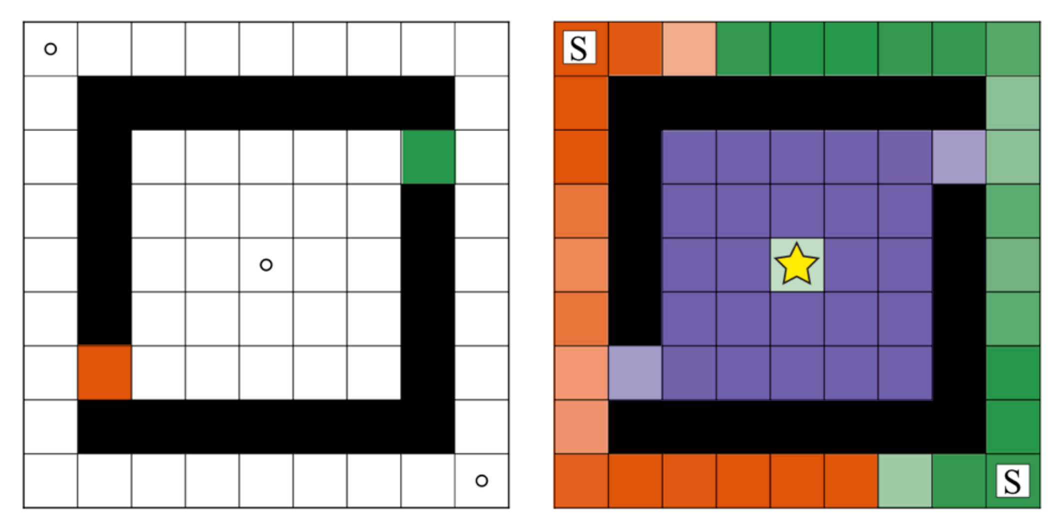

Solway et al. 2014 - Uninformed search

Results from Solway et al. 2014 (top row) and our replications (bottom row).

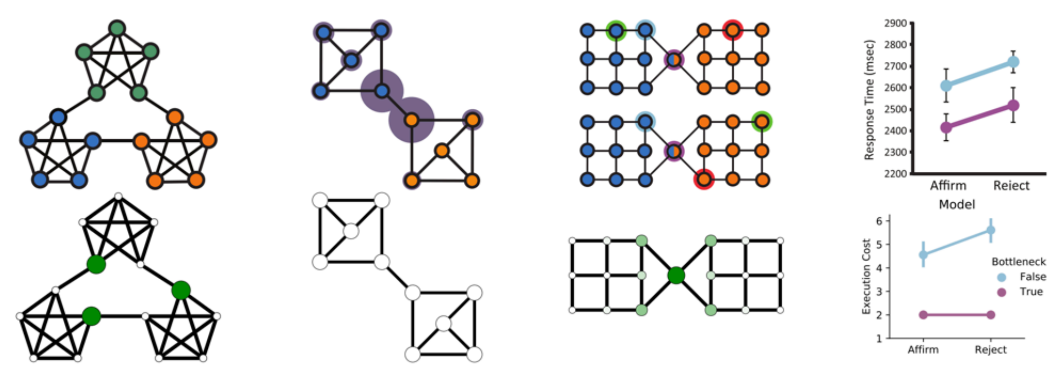

Solway et al. 2014 - Heuristic search

For structured tasks (in this case, Tower of Hanoi), we found that uninformed search alone did not identify bottleneck states.

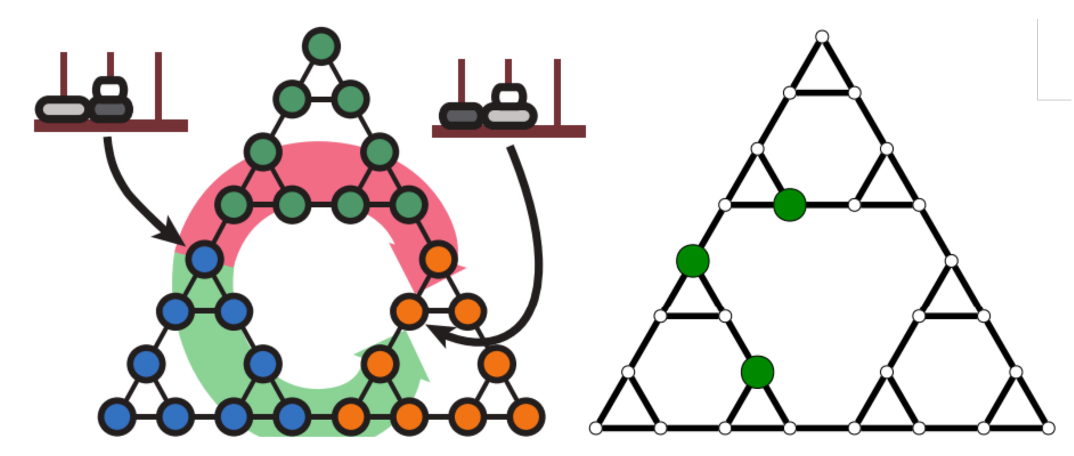

Encoding this structure in a heuristic used by search (left) was necessary to find subgoals related to bottlenecks (right).

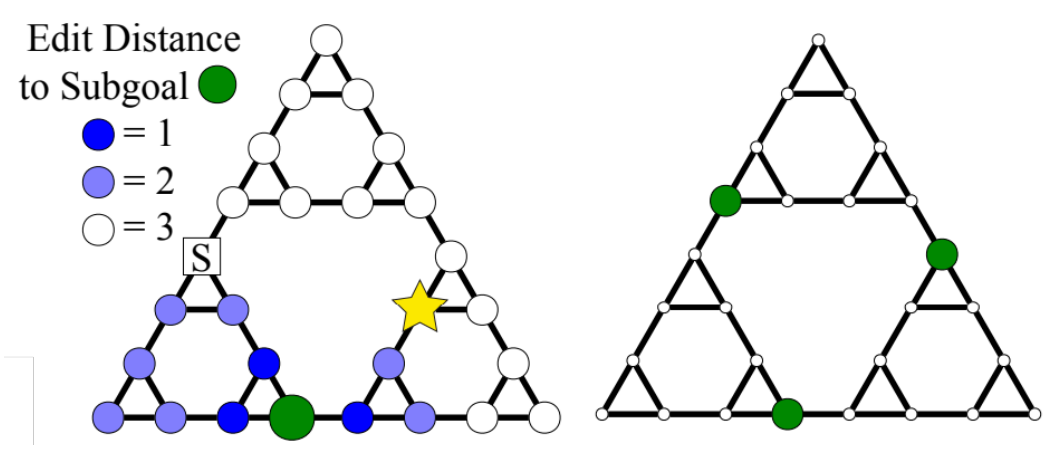

Conclusions and Future Work

- Previous work emphasizes task decomposition as inference about latent structure.
- Here, we formulate task decomposition as a resource-rational representation problem.
- Our model identifies intuitive subgoals that facilitate efficient planning computations and replicates previous experimental findings[4].
- Future work will empirically contrast task decomposition as latent structure inference versus resource-rational representation and further explore how planning algorithms shape task decomposition.

References

1. Bellman, R. (1957). Dynamic programming. Princeton University Press.
2. Kaplan, C. A., & Simon, H. A. (1990). In search of insight. Cognitive Psychology, 22(3), 374 - 419.
3. Sacerdoti, E. D. (1974). Planning in a hierarchy of abstraction spaces. Artificial Intelligence, 5(2), 115–135.
4. Solway, A., Diuk, C., Cordova, N., Yee, D., Barto, A. G., Niv, Y., & Botvinick, M. M. (2014). Optimal behavioral hierarchy. PLoS Computational Biology, 10(8), e1003779.
5. Tomov, M. S., Yagati, S., Kumar, A., Yang, W., & Gershman, S. J. (2020). Discovery of hierarchical representations for efficient planning. PLOS Computational Biology, 16(4).