

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Visual scoping operations for physical assembly

Permalink

<https://escholarship.org/uc/item/7cd0z0v7>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN

1069-7977

Authors

Binder, Felix Jedidja
Mattar, Marcelo G
Kirsh, David
et al.

Publication Date

2021

Peer reviewed

Visual scoping operations for physical assembly

Felix Binder

Dept. of Cognitive Science
UC San Diego
fbinder@ucsd.edu

Marcelo Mattar

Dept. of Cognitive Science
UC San Diego
mmattar@ucsd.edu

David Kirsh

Dept. of Cognitive Science
UC San Diego
kirsh@ucsd.edu

Judith Fan

Dept. of Psychology
UC San Diego
jefan@ucsd.edu

Abstract

Planning is hard. The use of subgoals can make planning more tractable, but selecting these subgoals is computationally costly. What algorithms might enable us to reap the benefits of planning using subgoals while minimizing the computational overhead of selecting them? We propose visual scoping, a strategy that interleaves planning and acting by alternately defining a spatial region as the next subgoal and selecting actions to achieve it. We evaluated our visual scoping algorithm on a variety of physical assembly problems against two baselines: planning all subgoals in advance and planning without subgoals. We found that visual scoping achieves comparable task performance to the subgoal planner while requiring only a fraction of the total computational cost. Together, these results contribute to our understanding of how humans might make efficient use of cognitive resources to solve complex planning problems.

Keywords: planning; problem solving; physical reasoning; spatial reasoning; task decomposition; hierarchical reinforcement learning

Introduction

Imagine you are preparing a meal. You need to wash, chop, heat, sauté, move around, boil. How do you know what to do next? Planning every action in advance is practically impossible: the number of potential sequences of actions grows exponentially with one’s action repertoire and the number of steps needed to reach the goal state. Not planning at all is no better: some actions must precede others (you can’t sauté without heating the pan), and timing matters too (e.g., adding sauce to uncooked pasta). How do humans manage to routinely solve such complex planning problems in everyday life?

Classical approaches to planning formulate such problems as search over a space of actions (Newell & Simon, 1972; Kirsh, 2009), augmented with heuristics and stochastic methods to more selectively search the space of possible actions (Geffner, 2013). Even with heuristics and stochastic methods, planning complex tasks in a rich environment is computationally prohibitive, because the space of potential plans that needs to be searched grows rapidly with both the number of possible actions available and the number of actions required to achieve the task (Bellman, 1957). The computational complexity predicted by the classical approach is at odds with the effortlessness with which people act in the real world (Kirsh, 2009).

A promising alternative approach from hierarchical reinforcement learning (Botvinick, Niv, & Barto, 2009) permits an agent to learn abstractions over sequences of actions, which can then be invoked as subgoals during plan-

ning (Maisto, Donnarumma, & Pezzulo, 2015; Zhang, Lerer, Sukhbaatar, Fergus, & Szlam, 2019; Bapst, Sanchez-Gonzalez, Shams, et al., 2019). However, while proposing good subgoals can reduce the computational cost of planning the sequences of actions (Correa, Ho, Callaway, & Griffiths, 2020), actually choosing which subgoals to propose can itself be highly costly. How might people manage these costs?

What these formal approaches tend to ignore is that when people confront such tasks they are often embedded in physical environments that can be reconfigured to suit their current goals (Kirsh, 1995). For example, one could gather all the vegetables on the cutting board and then focus on chopping what is there, ignoring what could be done elsewhere in the kitchen. After achieving this subgoal, it may be worth considering what the next subgoal should be (e.g., sautéing the vegetables), then focusing on that task (e.g., only considering actions available near the stove), and so on. We call this *visual scoping*—manipulating the visual environment to select the next subgoal. While such a strategy would not be expected to always identify the optimal sequence of subgoals, it may reduce the overall cost of jointly inferring subgoals and actions without leading to devastating consequences. This paper aims to establish a basic understanding of how such a visual scoping mechanism constrains planning behavior and impacts the overall computational cost of planning. Our investigation of visual scoping takes inspiration from recent work exploring how agents intervene on the world to aid in physical reasoning and planning (Dasgupta, Smith, Schulz, Tenenbaum, & Gershman, 2018; Allen, Smith, & Tenenbaum, 2020). Unlike these studies, we focus not on experimental or perceptual interventions which yield new information, but rather on interventions that select already present information to aid physical and spatial reasoning (Kirsh & Maglio, 1994).

As a case study, we consider planning in block-tower reconstruction problems, in which an agent uses an inventory of rectangular blocks to recreate a specific block tower. The combinatorial nature of construction leads to an explosion in the number of possible states to consider when planning block placements. To explore how visual scoping may influence problem solving in this challenging domain, we conducted a set of computational experiments comparing its behavior to that of both classical planning algorithms and more recent approaches which decompose the task into subgoals in advance of action-level planning (Correa et al., 2020). We discovered that visual scoping can approximate the success and




Figure 1: We relate three different strategies for planning using subgoals: no subgoals, planning only the next subgoal (scoping) and planning all subgoals in advance (full-subgoal planning). A schematic search tree of potential actions is shown at the bottom. Both the number of possible different actions (breadth of tree) and the length of sequences of action (depth) increases with the size of the subgoal.

efficiency of full task decomposition on a block-tower reconstruction task while requiring a much smaller computational budget overall. We also find that in trading off quick progress and low planning cost in choosing the next subgoals, valuing minimizing planning costs leads to more subgoals and poorer performance. Together, these findings help advance our knowledge of how perceptual and cognitive constraints interact to support efficient problem solving.

Approach

Block-tower reconstruction task

Three specific considerations motivated the choice of block-tower reconstruction for our experiments: (1) it requires planning over somewhat extended time horizons, (2) it is familiar enough to people that they can easily predict the consequence of their actions, and (3) it is inherently spatial, enabling visual scoping to select regions of the environment to focus on that could constitute plausible subgoals. Here we use the *block tower reconstruction task*, which entails assembly of two-dimensional block towers in a gridworld running simulated physics (McCarthy, Kirsh, & Fan, 2020). Similar block tower construction tasks have been used to study planning and physical reasoning in artificial agents (Sussman, 1975; Bapst, Sanchez-Gonzalez, Doersch, et al., 2019) and humans (Dietz, Landay, & Gweon, 2019; Cortesa, Jones, Hager, & Khudanpur, 2018; McCarthy et al., 2020).

On each trial, the planner is presented with an outline of a shape and has to recreate it by placing blocks from a fixed inventory in a building area. An action consists of the choice of a block (eg. 2×1) and a horizontal location. The block is then placed on top of the highest block or the ground in that horizontal location. Once a block has been placed it cannot be removed, necessitating planning in order to not get stuck

in dead ends.¹ As a simplification in our current experiments, blocks were also “glued down” after they were placed, preventing towers from toppling over. The trial ends either when the target shape is perfectly reconstructed, when no further block can be placed or when the planner chooses to not place a further block. The trial is considered successful if the target shape has been exactly reconstructed. 16 unique silhouettes were used in our experiments, spanning a range of difficulty levels for the planner. This is a challenging task for human participants: human subjects average 22.4% perfect reconstruction, albeit without glued blocks (McCarthy et al., 2020).

Visual scoping

The use of subgoals can reduce both the number of possible different actions and the length of sequences of actions needed to reach the (sub)goal (see Figure 1). Visual scoping attempts to reap this computational benefit while minimizing the cost of subgoal selection by only choosing the next subgoal depending on the current state of the environment. Our visual scoping planner works by interleaving two operations: (1) identifying the next subgoal to achieve, and (2) using an action-level search algorithm to plan the sequence of actions to achieve that immediate subgoal (Algorithm 1, Figure 1). Subgoals are defined spatially: a subgoal is defined as a region of space rather than as a specific state of the world. The set of potential subgoals we consider are horizontal layers, eg. the first or the first three rows of the target shape. Splitting a construction problem into horizontal layers is sensible given the incremental nature of building: it is easier to place blocks on top of other blocks rather than below them. Cortesa, Jones, Hager, Khudanpur, and Shelton (2017) show that young children naturally adopt a layerwise strategy on a related block construction task. In our block-tower reconstruction setting, this corresponds to identifying the height of the subtower the agent seeks to reconstruct, then searching over potential action plans for doing so.

Planners

Subgoal level planners We implemented two hierarchical subgoal planners, scoping and full-subgoal planning (see Algorithm 1 for a comparison of the algorithms), as well as a baseline of not using subgoals at all. The **scoping planner** considers all potential next subgoals given the current state of the environment. After a subgoal has been chosen, the subgoal is passed on to the action-level search algorithms and the actions it returns are applied to the environment. The process repeats until the target shape is completed or no subgoal that is perfectly solvable can be found. It has to trade off immediate progress and computational cost when selecting a subgoal: does it prefer more substantial subgoals (i.e., making

¹If actions can be undone, planning is not technically necessary (Kirsh, 2009): one could just try out actions and backtrack if they fail. However, placing and then undoing physical parts is costly and not always possible. Reasoning about the consequences of actions is thus important, particularly in domains where most paths do not lead to the goal.

Algorithm 1: Full-subgoal & scoping planner

Input: action-level search algorithm LLP , target shape G , environment E
 Parameters: weight parameter λ , computational budget b
 $c^{total} \leftarrow 0$

Run once:	<i>Full</i>
-----------	-------------

-4pt

While $E \neq G$:	<i>Scoping</i>
--------------------	----------------

$\Phi \leftarrow \{(g_{1:m1}), (g_{1:m2}), \dots, (g_{1:mn}) \text{ where } g_{m1} = g_{m2} = \dots = g_{mn} = G\}$ possible complete sequences of subgoals given E	<i>Full</i>
---	-------------

$\Phi \leftarrow \{(g_1), (g_2), \dots, (g_n) \text{ where } g_1, g_2, \dots, g_n > E\}$ possible sequences of subgoals given E of length 1	<i>Scoping</i>
---	----------------

For each (g) in Φ : $s = \text{Empty}$ For each g in (g) : $r_g \leftarrow \text{Area}(g) - \text{Area}(s)$ While $s_{final} \neq g$ and $c_g^{planning} < b$: $c^{solution} = 0$ $[(a_1, \dots, a_l)_g, c_g^{solution}, s_{final}] \leftarrow LLP(s, g)$ $c_g^{planning} += c_g^{solution}$ $c^{total} += c_g^{planning}$ if $s_{final} = g$: if subgoal is solved $s \leftarrow s_{final}$

else:

 $c_g^{planning} = \infty$ Continue to next $(g_k) \in \Phi$ For each (g) in Φ : $V_{(g)} \leftarrow \sum_{g \in (g)} r_g - \lambda * c_g^{planning}$

Apply $(a_1, \dots, a_l)_{g \in \text{argmax}(V_{(g)})}$ to E Choose sequence with highest score and apply all actions	<i>Full</i>
--	-------------

Apply $(a_1)_{g \in \text{argmax}(V_{(g)})}$ to E Choose sequence with highest score and apply only the first action	<i>Scoping</i>
--	----------------

define $LLP(s, g)$:Attempts to find a sequence of actions from s to g return $[(a_1, a_2, \dots, a_n), c^{solution}, s_{final}]$

rapid progress) even at the cost of higher computational cost, or more modest subgoals that are easier to solve? This trade-off is controlled by the λ parameter. The higher the value of λ , the more the scoping planner works to minimize action-planning cost. When $\lambda = 0$, the scoping planner maximizes progress no matter the cost. The **full-subgoal planner** finds one sequence of subgoals from the beginning of building to the final structure in one go. It first considers all possible sequences of subgoals that end in the full target shape, selects the one that minimizes the action-level search algorithm computational cost and builds the structure from start to finish, as in Correa et al. (2020). Finally, to compare the use of subgoal planners to a baseline of **pure action-level search** not using subgoals at all we also apply the pure action-level search algorithm directly on the target structure without decomposing

it into subgoals.

These planners use a given action-level search algorithm to determine for each potential subgoal the computational cost of solving the subgoal with the action-level search algorithm. This is done by running the action-level search algorithm repeatedly on a subgoal until a solution is found or the cost threshold b is exceeded. It is necessary to sample repeatedly, since the action-level search algorithms used here break ties randomly between equally good plans and therefore might yield different results. Since not all subgoals are achievable for a certain action-level search algorithm, a threshold is needed after which the subgoal is considered unsolvable. Only solvable subgoals are considered by the subgoal level planner. Finally, the subgoal planner selects the subgoal or sequence of subgoals that maximizes progress (the proportion of the target shape solved) while minimizing the cost of solving the subgoals using the action-level search algorithm. Because the subgoal planners minimize the planning cost of the action-level search algorithm, they perform resource-rational task decomposition. (Correa et al., 2020). Our implementation of resource-rational task decomposition differs from Correa et al. (2020) in three specific ways: (1) the use of abstract subgoal states (rather than specific world states serving as subgoals), (2) the ability to only plan a certain number of steps ahead as mediated through λ , and (3) the ability to handle impossible subgoals and stochastic action-level search algorithms.

Action-level search algorithms The subgoals planners are hierarchical agents: they find a subgoal decomposition both using a particular action-level search algorithm, and specifically for that particular algorithm. An action-level search algorithm searches the space of possible actions to find a sequence of actions to reach a certain goal. Since the space of potential states (different placements of blocks in the building area) is very large—there are roughly $35^{15} \approx 1.5 \times 10^{23}$ different states—exhaustively searching the entire space of states for a complete sequence of actions is not feasible. Therefore, the action-level search algorithms are implemented as lookahead planners: if a path to the solution cannot be found within a certain computational budget, the planner plans a sequence of actions as long as it can given a computational constraint with ties broken randomly, then takes the first action of that sequence and plans again, now on the basis of the state resulting from the action just taken. We chose two classical search algorithms: one performing brute search, the other performing search informed by a heuristic.

Breadth first search lookahead (BFS) exhaustively explores all possible placements of blocks in the target shape n steps into the future or until a perfect reconstruction is found, then chooses the sequence of actions that maximizes the area of the target shape filled out.

A* lookahead runs the A* search algorithm (Hart, Nilsson, & Raphael, 1968) either until a sequence of actions resulting in a perfect reconstruction is found or until a specified computational budget is exceeded. A* chooses to explore the

action that minimizes $f(s) = g(s) + h(s)$, where $g(s)$ is the cost of reaching the current state from the start measured in number of blocks placed down and $h(s)$ is the heuristic for how expensive the target is to reach from the current state. Here, $h(n)$ is the number of cells left in the target shape divided by the size of the smallest block, so it provides an upper bound for the number of blocks left to fill out the target shape.²

Hypotheses The full-subgoal planners considers the largest set of subgoal decompositions, so we expect that it has the highest rate of perfect reconstruction as well as the lowest action planning cost of the solutions themselves compared to the scoping planner. However, we hypothesize that this optimality comes at a steep cost: we expect the subgoal-planning cost incurred by the full-subgoal planner to be higher than that incurred by the other two. By contrast, we predict that the scoping planner will trade off higher performance and cheaper action-level planning for a large decrease in the cost of finding the subgoals. Finally, we predict that not using subgoals at all will lead to a much lower rate of perfect reconstruction and higher action-level algorithmic cost of the solution that is found.

Experiments We ran each planner 32 times on each of the 16 structures using different random seeds. For the scoping planner, we used 64 different values of λ each.

Results

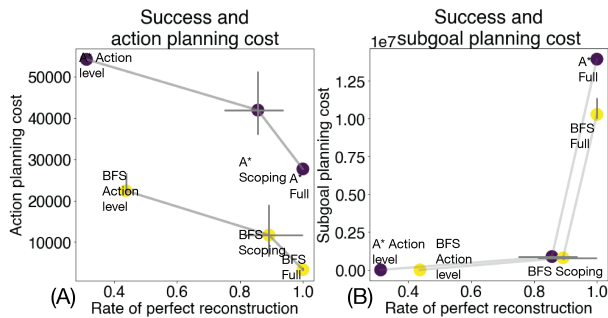


Figure 2: Scoping can approximate the success and action planning cost of full-subgoal planning while requiring a much smaller cost in choosing subgoals. (A) Each dot corresponds to a combination of subgoal planner (no subgoals, scoping, full-subgoal planning) and action-level search algorithm (A*, BFS). The cost of planning the actions given the sequence of subgoals and the rate of perfect reconstruction is shown. (B) Same as (A), but with the cost of selecting the subgoals shown instead. Costs are measured in evaluated number of states. Error bars represent bootstrapped 95% confidence intervals.

Pure action-level search To establish a baseline for both success and planning cost, the action-level search algorithms

²Note that the heuristic $h(s)$ is not strictly admissible: for states that cannot possibly reach a perfect reconstruction (for example, if a hole has been left and covered by a block), this locked in dead end is not going to be apparent in the heuristic.

Search	Budget	Accuracy	95% CI	Cost	95% CI
Random	None	0.038	[0.000, 0.125]	0	[0, 0]
BFS	Depth 1	0.100	[0.000, 0.250]	8.294×10^1	$[8.200 \times 10^1, 8.480 \times 10^1]$
BFS	Depth 2	0.187	[0.063, 0.313]	1.520×10^3	$[1.447 \times 10^3, 1.593 \times 10^3]$
BFS	Depth 3	0.263	[0.125, 0.438]	2.454×10^4	$[2.216 \times 10^4, 2.692 \times 10^4]$
BFS	Depth 4	0.410	[0.188, 0.625]	3.518×10^5	$[3.150 \times 10^5, 3.833 \times 10^5]$
A*	8 Iterations	0.062	[0.000, 0.188]	1.202×10^2	$[1.148 \times 10^2, 1.260 \times 10^2]$
A*	64 Iterations	0.083	[0.000, 0.188]	3.420×10^2	$[3.277 \times 10^2, 3.707 \times 10^2]$
A*	4096 Iterations	0.123	[0.000, 0.250]	1.653×10^4	$[1.453 \times 10^4, 1.853 \times 10^4]$
A*	65536 Iterations	0.100	[0.000, 0.250]	1.574×10^5	$[1.393 \times 10^5, 1.732 \times 10^5]$

Table 1: Increasing the budget of action-level search increases success somewhat, but costs grow exponentially. The table shows the cost and accuracy of various action-level search algorithms without resampling on a single attempt on the block construction task.

were used to try to find solutions to the 16 target shapes. Table 1 shows the success and planning cost of solutions found for a number of configurations of BFS lookahead and A* lookahead without resampling. Two particular action-level search algorithms will be used in conjunction with the scoping and full-subgoal planner: BFS lookahead with a search depth of 3 and A* lookahead with a budget of 4096 iterations. Using pure-action level search with resampling, BFS achieves a rate of perfect reconstruction of 0.438 (95% confidence interval (CI): [0.438, 0.438]; CI derived by computing the mean reconstruction rate across towers on each of 1000 iterations, where each iteration is defined by a novel permutation of the random seeds used to initialize search); A* achieves 0.313 (95% CI: [0.313, 0.313]). Block construction is a challenging task: pure action-level search is far from performing at ceiling. Increasing the search budget of the action-level search algorithms leads to an increase in success, but with it comes a ballooning of planning costs.

Full-subgoal planning We sought to replicate the finding that using subgoals in planning can reduce action planning cost (Correa et al., 2020). We apply the full-subgoal planner to the same 16 structures, using both BFS lookahead and A* lookahead as action-level search algorithms. The full-subgoal planner achieves a perfect reconstruction on every attempt. The action planning cost of those solutions is cheaper compared to pure action-level search (1.903×10^4 fewer states evaluated, 95% CI: $[1.790 \times 10^4, 2.354 \times 10^4]$, $p < 0.001$; A*: 2.651×10^4 , 95% CI: $[2.651 \times 10^4, 2.651 \times 10^4]$, $p < 0.001$; CI derived by computing the mean paired difference in planning cost between agents on each of 1000 iterations, where each iteration is defined by a new set of random seeds) compared to the respective action-level search algorithm without subgoals. While the chosen subgoals themselves are easy to solve, the cost of coming up with the sequence of subgoals is very large (BFS: $M = 1.028 \times 10^7$ states evaluated, 95% CI: $[1.001 \times 10^7, 1.139 \times 10^7]$; A*: 1.394×10^7 , 95% CI: $[1.394 \times 10^7, 1.394 \times 10^7]$).

Scoping To capture the planning behavior over a broad range of values for λ , we average over the dynamic range of λ (BFS: [0, 0.008], A*: [0, 0.003]) for the following anal-




Figure 3: An overview over the different decompositions chosen by each planner. Each vertical stripe corresponds to a single attempt of the planner to build a certain target shape. Each chosen subgoal is marked in a different color, with blue indicating a the first and orange the eight chosen subgoal. The columns are organized by target shape and for the scoping planners sorted by λ .

ysis. The scoping planner achieves a slightly lower success compared to full-subgoal decompositions perfect success rate (BFS: 0.104 lower rate of perfect reconstruction compared to full-subgoal planner, 95% CI: $[0.000, 0.188]$, $p = 0.042$; A*: 0.145, 95% CI: $[0.063, 0.250]$, $p = 0.008$). The scoping planner finds more expensive solutions (BFS: 9.107×10^3 more states evaluated, 95% CI: $[3.767 \times 10^3, 1.701 \times 10^4]$, $p < 0.001$; A*: 1.484×10^4 , 95% CI: $[8.268 \times 10^3, 2.546 \times 10^4]$, $p < 0.001$). However, the subgoal planning cost is dramatically lower (BFS: 9.528×10^6 fewer states evaluated, 95% CI: $[9.079 \times 10^6, 1.067 \times 10^7]$, $p < 0.001$; A*: 1.306×10^7 , 95% CI: $[1.299 \times 10^7, 1.316 \times 10^7]$, $p < 0.001$)—see Figure 2. Figure 3 shows the chosen task decompositions made by the planners.

When building block towers, the planner can choose to use large blocks, making progress quickly, or to apply a more conservative strategy of using smaller blocks. The average number of blocks in a successful solution measures this tendency. Using BFS lookahead, the scoping planner uses more blocks on average compared to pure action-level search (5.331 more blocks used per solution, 95% CI: $[2.688, 7.875]$, $p < 0.001$), but insignificantly fewer than the full-subgoal planner (2.272 fewer blocks, 95% CI: $[-1.689, 6.064]$, $p = 0.064$). Thus, the scoping planner using BFS lookahead finds somewhat longer solution solution compared to pure action-level search, but not full subgoal decomposition. Under A* lookahead, the scoping planner uses fewer blocks than both pure action-level search (1.414 fewer blocks, 95% CI: $[-2.500 \times 10^{-1}, 3.313]$, $p = 0.027$) and full-subgoal planning (3.095 fewer blocks, 95% CI: $[1.563 \times 10^{-3}, 6.188]$, $p = 0.015$).

Trading off immediate progress and action planning cost

To analyze the tradeoff between avoiding costs and making progress, we apply the scoping planner under a range of values for λ —see Figure 4. As expected, valuing reducing planning cost indeed results in a reduction in action planning cost (BFS: Pearson’s $r(62) = -0.927$, 95% CI: $[-0.929, -0.925]$, $p < 0.001$; A*: $r(62) = -0.581$, 95% CI: $[-0.581, -0.581]$, $p < 0.001$, all confidence intervals bootstrapped). Along with this comes a reduction in success (BFS: $r(62) = -0.955$, 95% CI: $[-0.955, -0.955]$, $p < 0.001$; A*: $r(62) = -0.840$, 95%

CI: $[-0.840, -0.840]$, $p < 0.001$): with lower action planning costs comes an increased chance of building oneself into a corner. When increasing λ , the number of subgoals also increases (BFS: $r(62) = 0.982$, 95% CI: $[0.976, 0.986]$, $p < 0.001$; A*: $r(62) = 0.964$, 95% CI: $[0.964, 0.964]$, $p < 0.001$). Paradoxically, trying to minimize action planning costs increases the total subgoal planning cost (BFS: $r(62) = 0.961$, 95% CI: $[0.950, 0.970]$, $p < 0.001$; A*: $r(62) = 0.920$, 95% CI: $[0.920, 0.920]$, $p < 0.001$). This is due to the increase in the number of subgoals, as smaller subgoals tend to be easier to plan. More subgoals means that the costly subgoal planning cost needs to be performed more often. This illustrates that cognitive resource limitations can lead to a larger total of cognitive resources used, albeit in smaller increments.

Finally, the found solutions themselves qualitatively differ across values of λ as well as action-level search algorithms. Increasing λ leads to more blocks being placed when using BFS as a action-level search algorithm, but a slight decrease with A* (BFS: $r(62) = 0.560$, 95% CI: $[0.420, 0.678]$, $p < 0.001$; A*: $r(62) = -0.054$, 95% CI: $[-0.248, 0.132]$ $p = 0.143$). Valuing making progress over avoiding planning cost thus leads the scoping planner with BFS lookahead to shorter solutions, whereas avoiding costs leads to more conservative, longer solutions.

Discussion

We found that both scoping and full-subgoal planning outperformed pure action-level search, consistent with the notion that decomposing complex problems can be beneficial. While full-subgoal planning succeeds in finding the best task decompositions, the cost of finding this decomposition into subgoals is much higher than the scoping planner requires to achieve comparable task performance. In order to be a plausible proposal for how people approach these planning problems, both kinds of costs need to be taken into account. Here we assume that subgoal planning cost as defined is a plausible proxy for the actual cost of finding subgoals: while humans likely don’t fully search for actions for all potential subgoals, the cost of subgoal planning arguably depends on the number of potential subgoals as well as their difficulty. Taken together, our findings suggest that visual scoping may be a promising way to reap the benefits of utilizing subgoals while minimizing the overhead of subgoal planning, thereby making efficient use of limited cognitive resources.

Given the way that scoping exploits spatial information to manage the computational overhead involved in planning, it may also be a useful source of insight into how people solve similar spatial reasoning problems. Towards this end, we are developing a novel behavioral paradigm allowing people to visually define subgoals during planning, enabling us both to make inferences about cognitive resource constraints and to observe how they interleave planning and action over time. A promising direction for future work is the evaluation of a broader array of scoping strategies, such as those forms of scoping closer to full-subgoal planning that plan more than

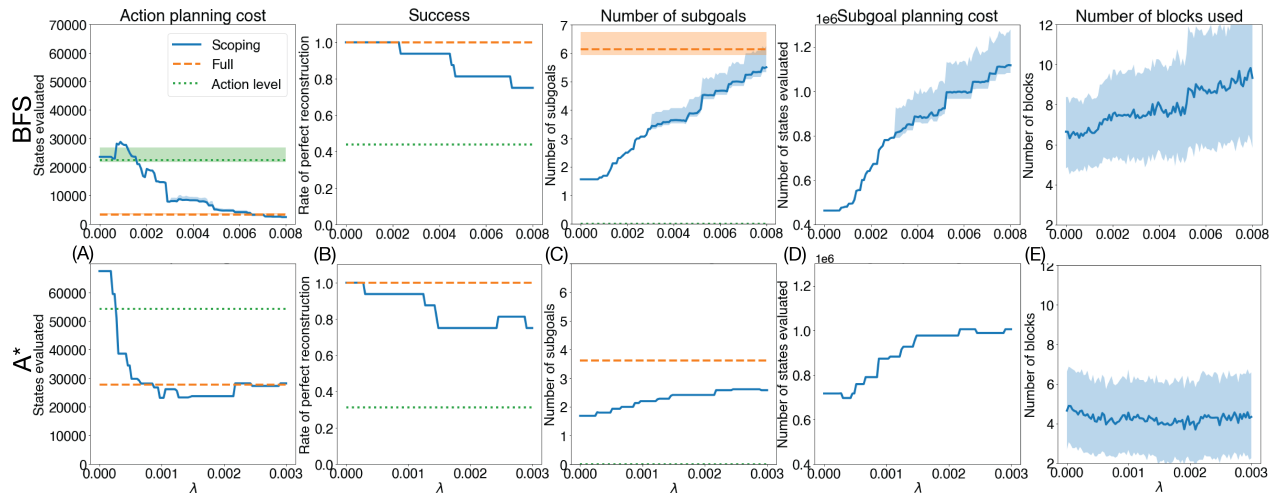


Figure 4: Parameter λ governs the tradeoff between choosing a subgoal that is easy to solve and one that maximizes progress. (A) A higher emphasis on minimizing action planning cost as opposed to maximizing progress (increasing λ) leads to a decrease in action planning cost, (B) a lower rate of perfect reconstruction, (C) an increase in the number of subgoals the planner actually ends up using, (D) as well as the total subgoal planning cost, since the scoping planners need to find the next subgoal more often. (E) It increases the number of blocks used for BFS lookahead, but not for A* lookahead.

one subgoal into the future at a time before taking action. Likewise, we aim to extend visual scoping from simple two-dimensional grid worlds to richer three-dimensional environments, where effective use of cognitive resources is crucial. In the block tower reconstruction task, the environment is deterministic. When the effect of the actions are uncertain—when the environment is stochastic or its dynamics aren’t fully known to the planner—the benefit of only planning subgoals into the near future is likely going to be more relevant. A future direction is to investigate how people scope in non-deterministic environments. Finally, visual scoping understands the planners as embedded in space and time: the planner exploits the structure of its environment to propose subgoals, and those subgoals are proposed on the basis of actions the planner has taken before. In this sense, visual scoping connects classical algorithmic models of problem solving with the notion of cognition being always situated in and dependent on an environment (Kirsh, 2009).

Acknowledgments

The authors would like to thank Will McCarthy for his development of the block tower reconstruction task, as well as members of the Cognitive Tools Lab at UC San Diego for their comments and support. This work was supported by NSF CAREER Award #2047191 to J.E.F.

All code and materials available at:
https://github.com/cogtoolslab/tools_block_construction

References

Allen, K. R., Smith, K. A., & Tenenbaum, J. B. (2020). Rapid trial-and-error learning with simulation supports flexible tool use and physical reasoning. *Proceedings of the National Academy of Sciences*, 117(47), 29302–29310.

Bapst, V., Sanchez-Gonzalez, A., Doersch, C., Stachenfeld, K. L., Kohli, P., Battaglia, P. W., & Hamrick, J. B. (2019). Structured agents for physical construction. *arXiv:1904.03177 [cs]*.

Bapst, V., Sanchez-Gonzalez, A., Shams, O., Stachenfeld, K., Battaglia, P. W., Singh, S., & Hamrick, J. B. (2019). Object-oriented state editing for HRL. *arXiv:1910.14361 [cs, stat]*.

Bellman, R. (1957). *Dynamic programming* (1st ed.). Princeton, NJ, USA: Princeton University Press.

Botvinick, M. M., Niv, Y., & Barto, A. G. (2009). Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective. *Cognition*, 113(3), 262–280.

Correa, C. G., Ho, M. K., Callaway, F., & Griffiths, T. L. (2020). Resource-rational Task Decomposition to Minimize Planning Costs. *arXiv:2007.13862 [cs]*.

Cortesa, C. S., Jones, J. D., Hager, G. D., & Khudanpur, S. (2018). Constraints and Development in Children’s Block Construction. *Proc. 40th Annu. Meet. Cogn. Sci. Soc.*, 6.

Cortesa, C. S., Jones, J. D., Hager, G. D., Khudanpur, S., & Shelton, A. L. (2017). Characterizing spatial construction processes: Toward computational tools to understand cognition. *Proc. 39th Annu. Meet. Cogn. Sci. Soc.*, 6.

Dasgupta, I., Smith, K. A., Schulz, E., Tenenbaum, J. B., & Gershman, S. J. (2018). *Learning to act by integrating mental simulations and physical experiments* (Tech. Rep.).

- Dietz, G., Landay, J. A., & Gweon, H. (2019). Building blocks of computational thinking: Young children's developing capacities for problem decomposition. *Proc. 41st Annu. Meet. Cogn. Sci. Soc.*, 7.
- Geffner, H. (2013). Computational models of planning. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4(4), 341–356.
- Hart, P., Nilsson, N., & Raphael, B. (1968). A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2), 100–107.
- Kirsh, D. (1995). The intelligent use of space. *Artificial intelligence*, 73(1-2), 31–68.
- Kirsh, D. (2009). Problem Solving and Situated Cognition. *The Cambridge Handbook of Situated Cognition*, 264–306.
- Kirsh, D., & Maglio, P. (1994). On Distinguishing Epistemic from Pragmatic Action. *Cognitive Science*, 18(4), 513–549.
- Maisto, D., Donnarumma, F., & Pezzulo, G. (2015). Divide et impera: Subgoalting reduces the complexity of probabilistic inference and problem solving. *Journal of The Royal Society Interface*, 12(104), 20141335.
- McCarthy, W., Kirsh, D., & Fan, J. (2020). Learning to build physical structures better over time. *Proc. 42nd Annu. Meet. Cogn. Sci. Soc.*
- Newell, A., & Simon, H. A. (1972). *Human problem solving* (Vol. 104). Prentice-Hall Englewood Cliffs, NJ.
- Sussman, G. J. (1975). *A computer model of skill acquisition*. New York; London: American Elsevier ; Elsevier.
- Zhang, A., Lerer, A., Sukhbaatar, S., Fergus, R., & Szlam, A. (2019). Composable Planning with Attributes. *arXiv:1803.00512 [cs]*.