A Task and Motion Approach to the Development of Planning

João Loula (jloula@mit.edu)
Department of Brain and Cognitive Sciences
Cambridge, MA 02139 USA

Kelsey R. Allen (krallen@mit.edu)
Department of Brain and Cognitive Sciences
Cambridge, MA 02139 USA

Joshua B. Tenenbaum (jbt@mit.edu)
Department of Brain and Cognitive Sciences
Cambridge, MA 02139 USA

Abstract

Developmental psychology presents us with a puzzle: though children are remarkably apt at planning their actions, they suffer from surprising yet consistent shortcomings. We argue that these patterns of triumph and failure can be broadly captured by the framework of task and motion planning, where plans are hybrid entities consisting of both a structured, symbolic skeleton and a continuous, low-level trajectory. As a proof of concept, we model two case studies from the tool use literature and show how their results can be understood by the interaction of symbolic and continuous plans.

Keywords: Planning; Tool use; Developmental Psychology; Action

Introduction

When a child is presented with a new problem to solve, where might she begin? We know a great deal about the content of children's intuitive theories about the world: from the earliest ages, they represent objects as discrete entities that cannot wink in and out of existence, and their physical understanding rapidly develops over the first two years of life. Simultaneously, we know that children are developing their motor capabilities over this period of time, including learning to walk, grasp, and perform simple tool manipulations. However, despite children's success in learning such a wide variety of skills, they show unexpected deficits in their ability to plan their actions. Most notably, children have consistently failed at a variety of tool use tasks, from bending a pipe cleaner into a hook shape to retrieve a toy from a tube, to stitching together two old plans to solve a new task. How is it possible that children can have rich intuitive physical theories, while not always being able to use these theories effectively to plan?

In this paper, we propose that an approach to the robotic planning in AI known as task and motion planning (Kaelbling & Lozano-Pérez, 2010) presents a grounded computational framework that can explain both childrens' striking successes and their striking shortcomings. In task and motion planning, simulations are guided by a symbolic search procedure that abstracts over continuous states. We argue that this view hints at new ways of looking at the development of planning: beyond just having more powerful simulators of the physical world, children might get better at planning in two distinct ways: by (a) creating better abstractions and (b) discovering better symbolic search procedures.

As a proof of concept for how task and motion planning can be used to explain the successes and failures of children's planning, we present two case studies of tool use from the developmental psychology literature that we model using this approach, and show that we can broadly capture developmental milestones.

Task and motion planning

Discrete planners, such as the General Problem Solver (Newell & Simon, 1961) or STRIPS (Fikes & Nilsson, 1971), are powerful because they can abstract over many possible states of the world. A simple symbolic expression like "pick up the rake" encapsulates many different states of an environment, and dramatically constrains the space of possible plans an agent has to consider. But this abstraction is also their weakness: a symbolic planner by itself cannot express *when* it is necessary to use a rake to reach an object, or *how* to use the rake. One might think of trying to incorporate information like this into a discrete planner, by adding expressions such as "if the object can't be reached, then use the rake". But this is sweeping the problem under the rug: to verify whether the object can or cannot be reached, one needs to grapple with the continuous geometric state of the world: where the object is placed, the shape of the surface it is on, whether there are obstacles preventing a reach, etc—one lesson learned from classical AI is that discrete planning problems quickly become intractable when trying to grapple with such information.

Continuous planners, on the other hand, deal beautifully with geometry, and can be used to solve many challenging motion planning tasks. They have no trouble generating continuous trajectories for grasping a rake, or for pulling an object with it; they are, however, incapable of deciding to use a rake in the first place, and unlikely to stumble open such a plan solely by moving through a sea of continuous trajectories.

The complementary strengths of discrete and continuous planners gave birth to task and motion planning (Kaelbling & Lozano-Pérez, 2010; Dantam et al., n.d.; Toussaint et al., 2018; Garrett et al., 2018), an approach where a robot makes a plan by first considering a discrete sequence of symbolic actions (a *task* plan) and then checking whether that leads to a feasible continuous trajectory (a *motion* plan). This approach has enjoyed great success in the robotics community, delivering impressive results on challenging, human-like sequential manipulation and tool use tasks.

This is distinct from alternative methods of planning: for example, adult decision making is primarily modeled using reinforcement learning, which tries to find ways of interacting with the world that lead to high reward (Daw et al., 2006; Niv, 2009). Reinforcement learning has been effective in explaining how adults approach tasks where they have no priors for how to interact with the scene, but cannot explain the flexibility people display when transferring knowledge from one task to another. If children are to become successful adult planners, they must learn how to plan in ways that generalize better than having stereotyped policies for each task they encounter.

Simply having a good model of the world is also not enough for planning. Naive approaches to model-based search would follow a "forward sampling" procedure in the space of low level motor actions (Betts, 1998). But searching for a plan in this space to accomplish any realistic goal that a child might have would be impossible – there would be far too many actions to consider. Instead, task and motion planning suggests searching through a more abstract, symbolic *task* space, which constrains the motor plans devised to reach a particular goal. If that is the case, then we should expect the development of planning in children to be driven not only by their proficiency in simulating and executing any one given course of action, but also by their proficiency in searching this symbolic space for relevant actions in the first place.

Planning Methodology

We operationalize task and motion planning by adapting the model introduced by Toussaint et al. (2018). Our goal is to make plans that manipulate many different objects in sequence and create and destroy multiple contacts: we do that by first sketching out how we want the geometric relationships between objects in our environment to evolve through time, and then attempting to fill that sketch with a trajectory that satisfies those relationships—the former is our task plan, the latter, our motion plan.

More concretely, given a set of geometric predicates a, we perform a breadth-first search through sequences of predicates $a_1, \ldots, a_K$, and for each sequence we attempt to find a trajectory $x_{0:T}$ that satisfies the constraints induced by those predicates at each timestep:

$$
\begin{aligned}
&\text{find } x_{0:T} \\
&\quad \text{s.t. } x_0 = s_0,\, c_{goal}(x_T) = 0 \\
&\forall t \in [0,T] : c_{path}(x_t, a_{k(t)}) = 0,
\end{aligned}
$$

where s_0 is the initial state, c_{goal} is the set of constraints that specify the goal, and $c_{path}(x_t, a)$ is the set of constraints encoded by a applied to the trajectory at time t.

The geometric constraints we use in this paper are meant to encode relationships between objects that are relevant to solving the task at hand, such as grasping a spoon, pulling an object with a hook, pushing it with a rod. Most task and motion planning approaches are silent on how an agent might learn about such constraints in the first place, and we shall not break with tradition here, focusing instead on how a given set of constraints can capture certain patterns of successes and failures in children's planning behavior.

Case Studies

We selected two studies from the literature on tool use in developmental psychology and recreated them in robotics simulations in order to have our model attempt to solve the same planning problems as the children in these studies did. The studies were picked for their striking patterns of failures and successes—in what follows, we summarize their results and present our versions of the tasks.

Case Study 1: Spoon grip

McCarty et al. (1999) introduce a task where a child between 9 and 19 months of age is presented with a spoon whose bowl is loaded with food, and then attempts to bring the food into their mouth by grasping and maneuvering the spoon. The child sits on their parent's lap, and the spoon placed upon a two-column support that crucially allows for it to be grasped from either above or below.

The design of the study hinges upon the fact that children of this age already have a dominant hand and exhibit a preference for using it. Trials were divided into *easy* and *difficult*: easy trials are the ones where the handle end of the spoon points to the same side as the child's dominant hand, difficult trials are the ones where it points to the opposite side. The reason for this distinction is that children prefer to grasp the spoon from above, and when doing that they'll be left with an adequate grip with their dominant hand on the easy trials, with their thumb pointing towards the bowl-end of the spoon, and an awkward grasp on the difficult trials, which makes it harder to get the food to their mouths (see the top part of figure 1 for an example of an awkward grasp on a difficult trial.)

There was a clear separation between how children different in age groups solved the task on difficult trials. 9-month olds mostly use the awkward grip, and placed the food into their mouths by the procedure shown in the top part of figure 1. 14-month olds also mostly used the awkward grip, but corrected it halfway through the trajectory by using many different strategies, including setting the spoon on the table in order to change the way they were gripping it. 19-month olds mostly chose the efficient grasp from the start.

Note that these differences cannot be due to how far ahead children are planning, as all of them grasp the spoon with the intention of ending the trajectory by placing it inside their mouths. The difference could, however, be due to the level of detail they are using to fill in this future trajectory, or the kind of abstraction they are deploying to represent it. Under that hypothesis, task and motion planning is a good candidate to model this phenomenon.

Experiment

We propose to model the developmental trajectory in the spoon task observed by McCarty et al. (1999) as a refinement of the kinds of abstractions children are using to plan for that task. More specifically, we propose that (a) 9-month olds' failures can be modelled by planning with a very coarse representation of the geometric relations which make for an effective grasp of the spoon (b) 19-month olds' successes can be modelled by a similar planning procedure, but with a more refined representation of what is required for an effective grasp, and (c) 14-month olds' mixed behavior can be modelled by plans that begin by using coarse heuristics similar to 9-month olds' and then replan with a more fine-grained model akin to 19-month olds'.

Environment: our version of the spoon environment can be seen in figure 1. We tried to maintain the core characteristics of the task in McCarty et al. (1999). The robot serves as a model for the child's dominant hand, as children in the task were for the most part insistent on only using that hand. A floating goal in red next to the robot stands for the mouth where the correct end of the spoon must be placed. The spoon, in its turn, rests in front of the robot upon a two-column support similar to the one in the original task: crucially, the spoon's bowl-end points away from the mouth, as in the difficult trials.

Task Predicates Used: only one task predicate was used per condition: the geometric constraint representing the gripper grasping the spoon. For the 9-month old model, this predicate depended only on the orientation of the gripper, so that the model would grasp the spoon the same way independent of how it was oriented. For the 19-month old model, the predicate enforced a particular alignment between the orientations of the spoon and of the gripper, akin to having one's thumb pointing towards the bowl-end of the spoon for a human hand. The 14-month old model was a mix of the two: in order to model the mid-trajectory replanning behavior seen in children this age, we took one of the intermediate states from the 9-month old model's trajectory and planned with the 19-month model starting from that state. The planners' tree search can also remove geometric predicates from previous timesteps—here this results in the gripper setting the spoon down and breaking the grasp like in figure 2.

Results

The 9-month old model's trajectory is shown in the bottom part of figure 1. Here, the naive geometric model of its grasp makes the robot, after the simple grasp on the second frame, have to settle for an awkward grasp to bring the spoon's bowl-end to its goal position.

The 14-month old model's trajectory can be seen in figure 2. Here, when asked to replan from a point in the middle of the 9-month old model's trajectory (frame 3), the model comes up with a solution for creating an efficient grip that involves setting the spoon down and letting go of it in order to change its grasp before bringing it to the destination.

The 19-month old model's trajectory, in contrast to the 9-month old's, goes for the grasp that would achieve the same relative orientation between gripper and spoon as an overhead grip in the easy condition—this results, in our model of the difficult condition, by it starting by grasping the spoon from below, which lands it in a comfortable grasp at the end of the trajectory, analogous to the 14-month old model.

Discussion

Our experiments offer an interpretation of the results in McCarty et al. (1999) where children learn more flexible abstractions for planning as they get older. Under that view, rather than just having a more powerful or more accurate simulator for guiding their actions, children would also develop by learning what kinds of abstractions over continuous states are safe to use in what situations, and —from their experience with spoons, 9-month olds would have come up with a rule that doesn't generalize well to the difficult trials in the study, whereas 19-month olds have learned a more general rule and also know when to deploy it.

The case of the 14-month olds is more interesting: though we recapitulated their behavior by forcing the model to replan, it's not clear why the children would choose the bad grasp and then switch mid-trajectory. If they're capable of understanding the utility of the correct grasp, why don't they deploy it from the start? One possibility is that, since the overhead grasp with the dominant hand is a simpler representation that is good enough for most situations they encounter, they plan with it by default, and that the meta-control that will allow them to correctly arbitrate between these representations won't develop until they get older.

Case Study 2: Sequential tool use

Metevier (2006) presented 36-month with two initial tasks, counterbalanced for order: one where they could retrieve an out of reach toy by pulling it with a rake, and another where they could retrieve a toy within a tube by pushing it out using a rod. These tasks were followed by a combination task, where once again the toy was inside a tube and could be pushed out with a rod, but now the rod was out of reach, and it was necessary for the child to first use the rake to retrieve the rod.

Though all children succeeded in both the rake and the rod task, only four out of sixteen of them managed to solve the combination task without assistance. After being given verbal hints to the solution, however—the strongest of which is "You can use this (the experimenter taps the head of the rake) to get this (tapping the rod) and then get the toy."—most children succeeded at the task.

Metevier (2006) writes of this pattern of failures: "Regardless of success rate, all of the children tested reached for the rod, asked for the rod, or stated that they needed the rod to solve the task. This suggests that the children did not initially understand that they needed to perform an intermediate step in the task that did not directly relate to getting the toy." It remains unclear, however, why that is: we know that children as young as 12-months old are apt at chaining actions together to achieve a goal (Sommerville & Woodward, 2005)—indeed, the very act of reaching for a tool like the children did in the

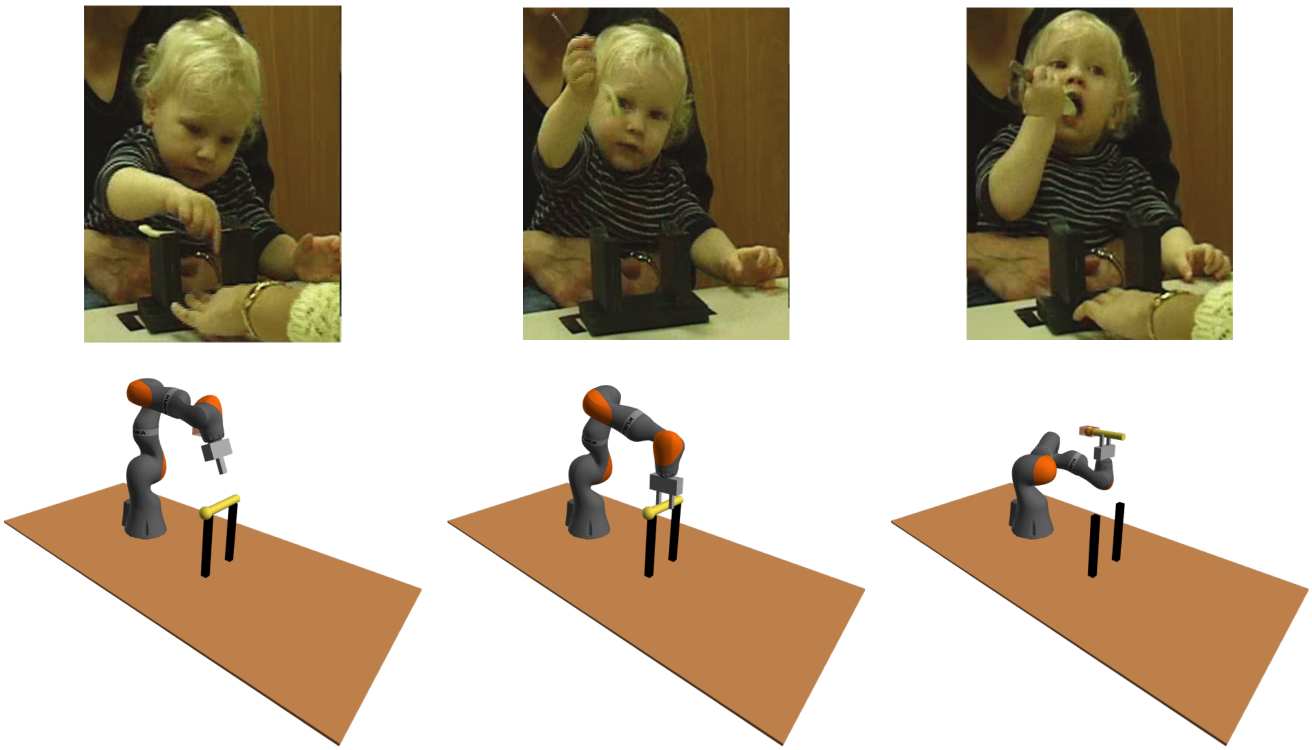

Figure 1: (a) a child performing the spoon task in study McCarty et al. (1999)—by choosing to grasp the spoon from above using their preferred hand, they finish the trajectory with an awkward grip (source: Keen et al. (2014).) (b) task and motion planning robot using coarse geometric primitives executing a similar grip in our version of the spoon task.

rake and the rod task is an example of performing an intermediate step not directly related to getting the toy; so it seems implausible that that in itself was the reason they struggled with the task.

Experiment

In the experiment that follows, we model the two initial tasks and the combination task using task and motion planning. We also model the verbal hint in the combination task as an alternative condition in which our model is provided with the correct task-level plan for the task but must discover the motion-level plan by itself.

Environment: three environments were part of this case study. The first environment is a model of the rod tasks, containing as objects a ball, a rod, and a tube; the second is a model of the rake task, containing a hook and a ball; the third one was the combination task, with a ball, a rod, a hook and a pipe (the third environment can be seen in figure 3). In all of the tasks, the goal is to get the blue ball, used as a stand-in for the toy in the original tasks, to a goal position, represented in red.

Task predicates used: for our task and motion planning model, we used a **slide** predicate, that the gripper, hook, and rod could apply to any of the objects (except for the pipe, which was fixed). The gripper was also capable of grasping both the hook and the stick, and any geometric constraint created during a plan could be destroyed—this happens for instance when setting down an object. For every condition, we restricted the planner's task predicates to only those relevant to the environment at hand. For any predicate and pair of objects, a pre-specified pair of points in their surface was used as the reference for the kinematic constraint.

Results

In table 1 we show the number of nodes searched and the CPU time elapsed until a solution was found in all four conditions,

In both the hook task and the rod task, our model searches through 3 nodes before finding a solution, which takes only a couple of seconds—after considering only grasping the tool, which doesn't help with the task at hand, or sliding the ball with the gripper, which is impossible due to it being out of reach, the model settles in both cases for using the tool to slide the ball.

The combination task is dramatically more difficult, taking an order of magnitude more time to solve and searching through two orders of magnitude more nodes. This happens because the string of actions required for the final plan is extremely long: picking up the hook, pulling the rod, releasing the rod, setting down the hook, picking up the rod, and finally pushing the ball: see figure 3 for the final trajectory found by our planner. When doing symbolic search, the number of nodes increases exponentially with the depth of the tree, which is the reason for the more than a hundred nodes searched—the corresponding increase in time is less dramatic

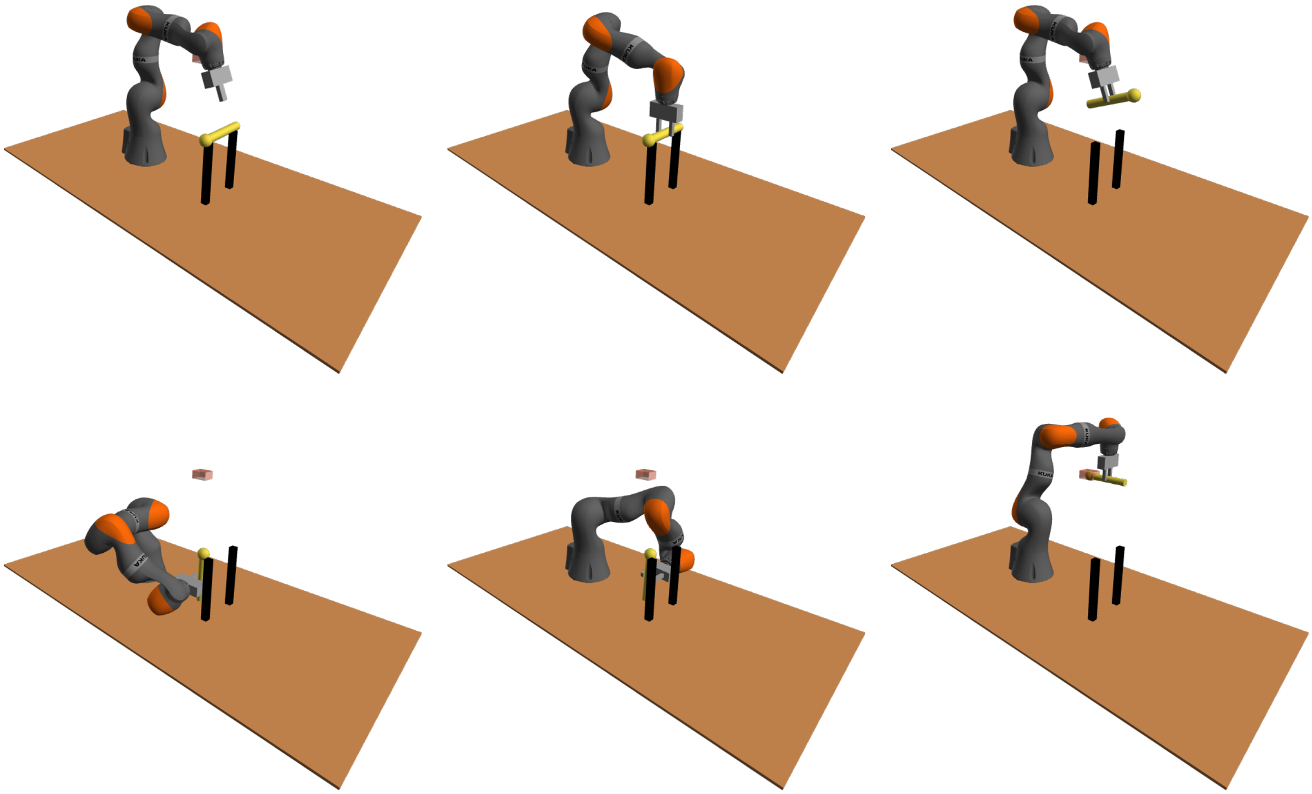

Figure 2: Task and motion planning model of 14-month olds in the spoon task.

in our model because the brunt of the search time is spent on a few nodes that cannot be quickly discarded by the continuous solver.

On the other hand, if the model receives a hint in the form of the ground truth task plan to be transformed into a motion plan, the problem of doing task-level search is removed and the plan is found as efficiently as those in the single tool conditions.

Table 1: Number of task nodes searched and duration of planning in CPU time for the four conditions in case study 2.

Task	nodes searched	CPU time (s)
Pull with hook	3	1.67
Push with rod	3	3.31
Combination	105	27.98
Combination (with hint)	1	1.77

Discussion

These experiments offer a simple interpretation of the results found by Metevier (2006): even though children are likely to reason symbolically about tool-use and use that reasoning to successfully guide their low-level plans when retrieving a toy with a rake or a rod, the apparent simplicity of a plan that requires chaining these two actions together can actually render the problem intractable for some types of planners by blowing up the search space.

Under this view, the problem is not with performing actions not directly related to the goal, which children successfully do in this and other tasks, but in discovering procedures to more efficiently navigate such large search spaces. The blow-up we observe in the combination task stems mostly from the choice of primitives and the naive breadth-first search procedure for the task-level plan—it is likely that, as children grow older, they develop sophisticated abstractions and heuristics for planning, avoiding such tractability problems. For instance, one could imagine that after solving the first two tasks, a child could learn to represent each of them as a single, more abstract predicate—in that case, the combination task would be significantly easier, requiring only two rather than six predicates to solve.

A crucial aspect of the original study is the fact that most children succeeded at the task after receiving a verbal hint describing both the correct sequence of tools to use and on which objects. This is important because it squarely places the 3-year olds' initial failure at this task as a problem of search: they are perfectly able to represent and execute the solution after being directed towards it, but just fail to find that solution on their own. One of the strengths of task and motion planning is providing an explanation for how such linguistic input might help one search through a space of con-

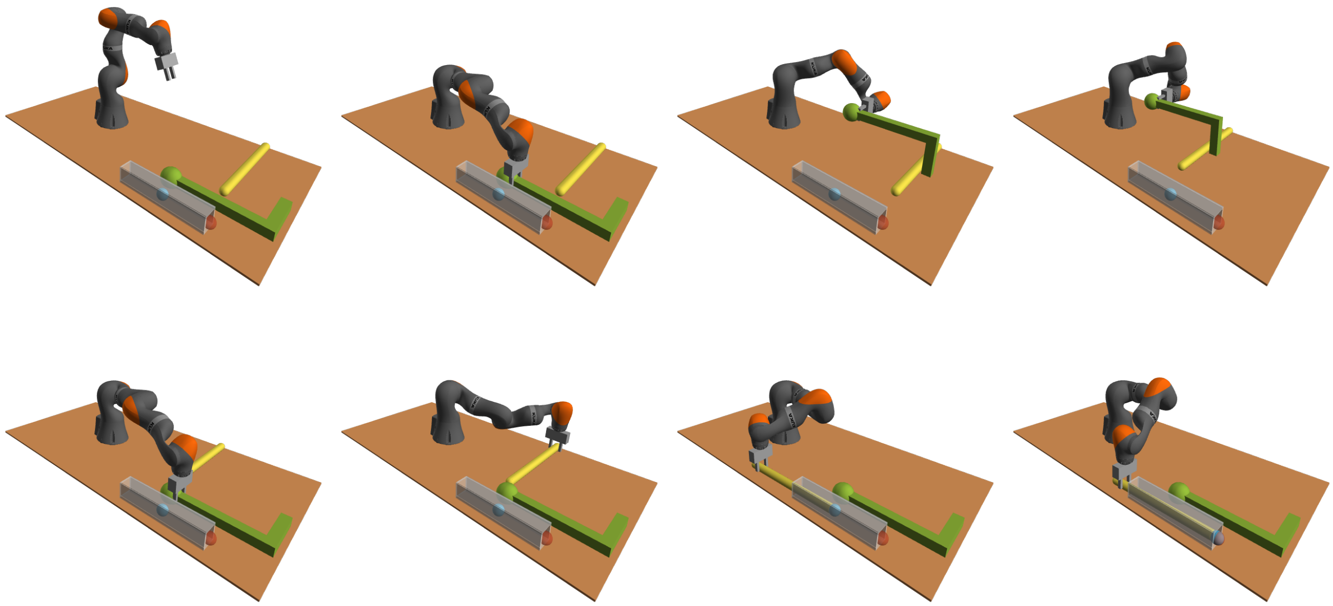

Figure 3: Discovered solution the combination task in case study 2. The goal is to get the blue ball to the position in red.

tinuous motor actions—most approaches to planning operate solely in the latter domain and would be unable to make use of such input.

We want to stress that this line of reasoning is similar in spirit to how the result was interpreted by the original authors: Keen (2011) writes “It is not clear how the children used this hint to guide their action, but one possibility is they visualized themselves making the sequence of actions. If so, they could subsequently carry out these actions in the order indicated and achieve success.” We see the contribution of our work as formalizing what it means for a simulation to be guided by a symbolic representation of a plan.

References

Betts, J. T. (1998). Survey of numerical methods for trajectory optimization. *Journal of guidance, control, and dynamics*, *21*(2), 193–207.

Dantam, N. T., Kingston, Z. K., Chaudhuri, S., & Kavraki, L. E. (n.d.). Incremental task and motion planning: A constraint-based approach..

Daw, N. D., O’doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, *441*(7095), 876–879.

Fikes, R. E., & Nilsson, N. J. (1971). Strips: A new approach to the application of theorem proving to problem solving. *Artificial intelligence*, *2*(3-4), 189–208.

Garrett, C. R., Lozano-Pérez, T., & Kaelbling, L. P. (2018). Sampling-based methods for factored task and motion planning. *The International Journal of Robotics Research*, *37*(13-14), 1796–1825.

Kaelbling, L. P., & Lozano-Pérez, T. (2010). Hierarchical planning in the now. In *Workshops at the twenty-fourth aaai conference on artificial intelligence.*

Keen, R. (2011). The development of problem solving in young children: A critical cognitive skill. *Annual review of psychology*, *62*, 1–21.

Keen, R., Lee, M.-H., & Adolph, K. (2014). Planning an action: A developmental progression in tool use. *Ecological Psychology*, *26*(1-2), 98–108.

McCarty, M. E., Clifton, R. K., & Collard, R. R. (1999). Problem solving in infancy: the emergence of an action plan. *Developmental psychology*, *35*(4), 1091.

Metevier, C. M. (2006). *Tool-using in rhesus monkeys and 36-month-old children: acquisition, comprehension, and individual differences*. University of Massachusetts Amherst.

Newell, A., & Simon, H. A. (1961). *Gps, a program that simulates human thought* (Tech. Rep.). RAND CORP SANTA MONICA CALIF.

Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, *53*(3), 139–154.

Sommerville, J. A., & Woodward, A. L. (2005). Pulling out the intentional structure of action: the relation between action processing and action production in infancy. *Cognition*, *95*(1), 1–30.

Toussaint, M., Allen, K., Smith, K. A., & Tenenbaum, J. B. (2018). Differentiable physics and stable modes for tool-use and manipulation planning. In *Robotics: Science and systems.*