

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Building blocks of computational thinking:Young children's developing capacities for problem decomposition

Permalink

<https://escholarship.org/uc/item/5tt891kg>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 41(0)

Authors

Dietz, Griffin

Landay, James A.

Gweon, Hyowon

Publication Date

2019

Peer reviewed

Building blocks of computational thinking: Young children’s developing capacities for problem decomposition

Griffin Dietz¹, James A. Landay¹, Hyowon Gweon²

{dietz, landay, hyo}@stanford.edu

¹Computer Science Department, Stanford University

²Department of Psychology, Stanford University

Abstract

Computational thinking (CT) refers to a range of problem-solving skills applicable to computer science and everyday life. Although recent research in developmental cognitive science suggests mental capacities relevant to CT may emerge quite early in life, research on CT, and computer science education more generally, has made little contact with this literature. As a way to better bridge these fields, we explore the development of problem decomposition, a critical feature of CT, in the spatial domain. We ask whether young children can break a complex spatial problem down into subcomponents that can be reassembled to solve the overarching problem. Across two experiments (Exp.1: 4- to 7-year-olds; Exp.2: 3- to 5-year-olds) that involve constructing block structures, we demonstrate that some of the key capacities underlying problem decomposition begin to emerge in preschool years and develop throughout early childhood. Although preschool-aged children struggle to solve an open-ended decomposition problem that requires generation and execution of decomposition plans, even 4-year-olds can successfully evaluate the viability of these plans. These results suggest that experimental methods in developmental cognitive science can inform CS education research that focuses on promoting CT; by identifying when and how CT concepts emerge in early childhood, we can better create age-appropriate educational tools.

Keywords: computational thinking; problem decomposition; problem solving; cognitive development; intuitive physics

Introduction

The ability to break down a large problem into smaller parts is important for many real-world tasks. To decompose a problem effectively, one must understand its constraints, generate potential solutions, and evaluate the strengths and weaknesses of those solutions. Importantly, these steps are often better taken *before* one actually acts; attempts to achieve a complex task without proper planning can lead to unnecessary effort to correct a mistake or even irreversible failure. But what does it take to be good at problem-solving and planning?

More than a decade ago, Wing (2006) popularized the concept of computational thinking (henceforth CT). CT is a term that collectively refers to a range of skills that are crucial to effective problem-solving, and it incorporates various cognitive strategies considered fundamental to computer science (CS) (Wing, 2006; Barr & Stephenson, 2011; Brennan & Resnick, 2012). Mental activities like abstraction (i.e., generalizing problem features to preserve only relevant information; Kramer, 2007) and problem decomposition (i.e., breaking a complex problem into solvable subcomponents; Barr & Stephenson, 2011) are key components of CT. Indeed, these skills are critical to building good computer programs; anyone who has engaged in programming understands the importance of abstracting away from a problem to identify its basic

structure and decomposing that structure into solvable parts.

Yet, the importance of CT reaches far beyond programming (Wing, 2006). Abstract thinking, problem decomposition, and the ability to evaluate potential plans are skills that allow us to tackle a range of everyday tasks as well as larger, more complex problems that involve multiple sub-goals, such as conducting scientific research or building an architectural project. In particular, to successfully achieve these larger goals, one must: (1) represent the current state of the world (i.e., what does the empty lot look like?, what materials do we have?) as well as the state of the desired end-goal (i.e., what do I want to build?), (2) identify the units that comprise the end goal (i.e., what sub-goals should I complete?) and construct the possible future states from applying these units (i.e., what will the structure look like given these components?), and (3) evaluate the viability and effectiveness of different sets of potential units and interventions (i.e., should we build the columns or the roof first?, which size columns are most suitable?). In other words, effective problem-solving involves the representational and inferential abilities to *generate* possible ways to decompose the problem space and *evaluate* the viability of a potential decomposition plan. By engaging in these mental processes prior to executing a given plan, one can solve a problem with less trial-and-error.

While CT has been a useful construct to raise awareness of the relevance of these skills in both computing and everyday life, it remains a difficult concept to operationalize or measure. This difficulty may arise from the fact that CT is not a single thing; it is a collection of various mental operations whose cognitive mechanisms are poorly understood. Furthermore, although CT presumably involves reasoning abilities that have been topics of interest in cognitive development research, this body of work has remained rather disconnected from the literature in CS education, leading many CS educators to believe that CT develops relatively late in childhood (Guzdial, 2015). Our goal is to take a step towards synthesizing these fields, and build on prior work to ask whether the ability to decompose a complex problem—a key component of CT—is present early in life. In the following, we summarize related work on young children’s inferential capacities and introduce a novel task for testing problem decomposition.

Prior work in cognitive development has revealed rich, sophisticated abilities in young children to engage in abstract reasoning and learning (Gopnik, 2012; Schulz, 2012). Although these studies are primarily aimed at identifying the developmental origins of the human ability to engage in sci-

entific thinking, collectively their findings suggest that the basic representational and inferential capacities supporting CT may emerge much earlier than previously thought. For instance, preschool-aged children construct novel hypotheses from observations via inductive generalization and design novel experiments to test these hypotheses by engaging in selection and isolation of relevant variables (Cook, Goodman, & Schulz, 2011; Legare, 2012). These abilities are foundational to successful problem solving.

Furthermore, CT involves the understanding that good plans achieve a goal effectively and efficiently. Evidence suggests the rapid development of planning abilities between ages 4 and 6, including an increase in the number of steps children can plan ahead to solve a problem (Klahr & Robinson, 1981) and improvements in the ability to deploy appropriate strategies depending on the task (Gardner & Rogoff, 1990). Prior work has also shown that even infants expect rational agents to act in ways that minimize cost (Gergely, Nádasdy, Csibra, & Bíró, 1995; Scott & Baillargeon, 2013), and they infer the reward an agent assigns to a goal based on the cost incurred to achieve it (Liu, Ullman, Tenenbaum, & Spelke, 2017). By age 5, children can even design informative experiments to infer the subjective costs or rewards of achieving a goal (i.e., an agent's competence or preferences) by systematically manipulating the objective rewards or costs of completing a task (Jara-Ettinger, Gweon, Tenenbaum, & Schulz, 2015; see Jara-Ettinger, Gweon, Schulz, & Tenenbaum, 2016 for a review). Collectively, these early-emerging capacities to generate and test hypotheses, engage in advance planning, and reason about efficiency suggest that the basic aspects of CT may emerge earlier than commonly believed.

Building on this prior literature, we designed a novel block-building task to study one of the key components of CT: problem decomposition. Block-building tasks are familiar to young children, and have historically been considered a useful domain for studying the development of planning and problem-solving. Block construction has been shown to be an indicator of early spatial skills (e.g., mental rotation, Brosnan, 1998), which correlate highly with later success in programming and STEM (Cooper, Wang, Israni, & Sorby, 2015; Verdine et al., 2014; Wai, Lubinski, & Benbow, 2009). Thus, studying children's ability to generate and execute an effective block-building plan can provide a unique window into understanding the early development of CT. Yet, prior work is largely limited to exploring children's bottom-up building processes, allowing them to build the target structure in a piecemeal manner. Whether or not children can engage in top-down problem decomposition remains an open question.

A key strength of our task is that it requires more than merely copying a model block structure: children must figure out a viable plan within the constraints of the task by decomposing the structure into appropriate parts. In simple block-building tasks, one might succeed by accumulating raw materials (i.e., individual blocks) in a piecemeal fashion. Similar to the ways beavers or birds build their dams or nests, a

child could repeatedly stack blocks to create a tower. However, imagine a child wants to build a structure resembling the bridge in Figure 1. Simply accumulating individual blocks isn't sufficient; the child must first assemble the "legs," and then place a horizontal bar on top. If a child starts by creating "pillars" that are as tall as the bridge itself (3 blocks), then a single block in the middle would not stay in place. This example demonstrates how a bottom-up building process can be insufficient even for seemingly simple tasks. Rather, this problem resembles the way that we approach larger, real-world engineering projects; we must take the desired goal, break it into smaller problems, and determine how those components should be solved and assembled within the constraints of the task. Thus, the goal in designing our task was to provide a context in which children would approach a complex problem in a similar manner under clearly defined task parameters.

Recent work demonstrates that both adults and children leverage intuitive physics when evaluating the stability of block structures (Battaglia, Hamrick, & Tenenbaum, 2013; Kamps et al., 2017; Yildirim, Gerstenberg, Saeed, Toussein, & Tenenbaum, 2017), that children as young as 4 can gauge the difficulty of building such structures (Gweon, Asaba, & Bennett-Pierre, 2017), and that they are capable of copying a model block structure when given the required pieces (Cortesa et al., 2018). Critically however, start-to-end construction requires intelligently *generating* those pieces as well. Computational models optimized to generate instructions for the construction of block structures identify structural components while accounting for the effect of gravity on future layers (Zhang, Igarashi, Kanamori, & Mitani, 2017). However, the ability to determine the required subcomponents based on an intuitive understanding of task-specific constraints has not been tested in young children, even though such ability might provide the key foundation for a more general ability to engage in problem decomposition.

In Experiment 1, we embedded the process of *generating*, *evaluating*, and *executing* an appropriate decomposition plan into a fun, engaging block-building task. Given a target structure, children had to identify the underlying substructures, simulate ahead to determine if those substructures could combine into a self-supporting building, and then execute this plan to complete the task. In Experiment 2, we use a simplified version of the task to ask whether young children's difficulty in Experiment 1 comes from the process of *generating* a plan with the appropriate subcomponents, rather than the process of *evaluating* the viability of a given plan by engaging in physical simulation.

Experiment 1

Methods

Participants A total of 112 children (Age: 4.00–7.99) were recruited from a local children's museum and a university-affiliated preschool (38 4-year-olds: $M = 4.56$ (4.04–4.99); 31 5-year-olds: $M = 5.43$ (5.01–5.96); 23 6-year-olds: $M = 6.49$ (6.14–6.99); 20 7-year-olds: $M = 7.61$ (7.08–7.99)). They

were randomly assigned to either the Standing Bridge condition (N=62) or the Sideways Bridge condition (N=50). We planned to recruit at least 40 children in each condition who successfully completed the task (10-12 in each age group). Twenty-nine children were unable to accomplish the task, and the successful subgroup included N=42 in the Standing Bridge condition and N=41 in the Sideways Bridge condition. An additional 20 participants were dropped from analyses because they: (1) did not speak English (N=3), (2) ended the study early (N=7), (3) failed the warm-up task (N=6, see Procedure), or (4) the experimenter made an error (N=4).

Stimuli For the main test trial, the model bridge was comprised of seven one-inch wooden cubes (three across the top and two on each side as supports) painted metallic silver. In the Standing Bridge condition, this bridge was presented upright, and in the Sideways Bridge condition, it was lying down (see Fig. 1). Children were given 3 individual unpainted wooden blocks with which they could create substructures using a “magic box.” The magic box (see Fig. 1) was a cardboard box covered in felt with a small coin slot and an output window. A wire connected this box to the “construction zone” (a flat piece of foam core covered in black felt) and a large plastic button, suggesting they were all part of one causal mechanism. A coin was required each time the child operated the magic box to create a substructure. Inside the box were pre-assembled metallic structures. For the main task, where children had 3 individual blocks, there were 4 possible configurations of shapes that children could make; we prepared 4 metallic structures for each shape for a total of 16. We included additional structures for the practice trials.

Procedure Children were introduced to the experimenter’s “magic box” which could turn a set of one-inch wooden blocks into larger, metallic blocks of varying shapes. The child had to first build a structure on the construction zone using individual wooden blocks; after putting a coin in the slot, children could press the plastic button to generate a single metallic block that had the same shape as the structure on the construction zone. In reality, when the child operated the magic box, the experimenter surreptitiously reached into

the box through a hidden opening, found the corresponding metallic block, and placed it in the magic box output window (see Fig. 1). From the child’s view, this created the illusion that the child’s button press operated the magic box to generate the metallic structures.

A brief warm-up task ensured that the child understood the purpose of the main task and how to operate the magic box. In the warm-up, the child was given four wooden blocks, and was asked to build a 4-block ‘T’ shape (Trial A), a 3-block ‘L’ shape (Trial B), a single block (Trial C), and an 8-block cube (2x2x2) composed of two 4-block squares (Trial D). We excluded children who failed to complete this pretrial task from subsequent analysis to ensure that all children included in the study understood the magic box paradigm and were able to use the magic box to build metallic block pieces.

In the main task, children were asked to build a bridge using the blocks and the magic box. Critically, children had only three wooden blocks such that the metallic blocks they could build using the magic box was limited to a particular set of shapes. They were also given 11 coins; this limited the number of possible times children could generate a metallic structure, providing a pressure to solve the task efficiently.

In the Standing Bridge condition, the upright bridge was subject to the forces of gravity, and thus required a specific block set and assembly sequence (i.e., set up two 2-block bases and place a 3-block horizontal bar on top). The Sideways Bridge condition was identical to the Standing Bridge condition except that the bridge was laid flat (and thus not subject to the force of gravity). While the task still forced children to decompose the structure, there were multiple possible solutions and the order of construction did not matter. Thus, the Sideways Bridge condition still required the ability to follow task instructions, create parts, and assemble the final structure. However, the need to engage in advance planning to generate the “correct” decomposition plan and evaluate its viability was not as critical for success.

The children were given up to 10 minutes to build the bridge, after which the experimenter stepped in to help and ended the study. Often the child got stuck (signaled by asking for help or a period of inactivity) or distracted, so to encourage the child to reengage, the experimenter offered one of two pre-scripted prompts. Additionally, after an extended period of inactivity or running out of coins, children were given the option of restarting the task in the remaining time.

Results and Discussion

This was an exploratory study to see whether children could engage in effective problem decomposition, rather than a test of a priori hypotheses. However, we could imagine seeing a few general trends in the data. First, we expected that children would become more successful and more efficient at completing the task with age. We measured efficiency using two different metrics: completion time (in seconds) and number of coins used (3 was the minimum). Second, independent of increasing performance with age, we also expected that children would perform better (i.e., higher success rate, as well




Figure 1: Left: Experimental setup. A) magic box, B) coin slot, C) output window, D) 1” wooden blocks, E) construction zone, F) plastic button; Right: Standing bridge (top) and sideways bridge (bottom).




Figure 2: (A) Success rate in each age bin and in each condition. (B-C) Time-to-completion (B) and number of coins used to complete the task (C) among the successful children, split by age and condition. Error bars: bootstrapped 95% CI.

as more efficient completion) in the Sideways Bridge condition than the Standing Bridge condition, because the Standing Bridge had just one viable 3-part decomposition solution. Third, we predicted an age by condition interaction; the condition difference would decrease with age as children become more proficient at finding solutions without trial and error.

A logistic regression with condition (discrete) and age (continuous) confirmed a relationship between age and success rate ($z = 2.41$, $p = 0.02$). However, we did not see a significant effect of condition ($z = 0.47$, $p = 0.64$) nor an interaction between age and condition ($z = -0.73$, $p = 0.47$).

Given the increase in success rate with age, we further analyzed data from the 83 successful children to see if children become more efficient problem-solvers with age. First, we looked at time-to-completion; a linear regression with both age (continuous) and condition (discrete) as predictors showed that older children take a shorter amount of time to complete the task ($t = -4.47$, $p < .001$). While children in the Sideways Bridge condition did not complete the task faster than those in the Standing Bridge condition ($t = -1.73$, $p = 0.09$), we did find an interaction between age and condition. However, the effect was in the opposite direction than we had initially predicted: the difference in completion time between conditions increased with age ($t = 2.06$, $p = 0.04$).

Another measure of efficiency—the total number of coins used—also showed a similar pattern. A linear regression on the 83 children who successfully completed the task revealed that the number of coins children used to complete the building task decreased with age ($t = -3.14$, $p = .002$); also consistent with time-to-completion, we did not find an effect of condition ($t = -1.49$, $p = .14$) but the difference between conditions increased with age ($t = 1.96$, $p = .05$).

We then looked at the proportion of children who completed the task with maximal efficiency (i.e., successfully building the bridge using just 3 coins). A logistic regression with condition and age showed an effect of age ($z = 3.61$, $p < .001$). Children were also more likely to perform optimally in the Sideways Bridge than in the Standing Bridge condition ($z = 2.08$, $p = 0.04$), and this tendency increased with age (age by condition interaction, $z = -2.54$, $p = 0.01$).

Overall, data from this exploratory study showed a few notable patterns. First, unsurprisingly, children became more successful and more efficient at solving the task with age across a number of measures: success rate, time-to-completion, and number of coins used to finish the task. These results are consistent with prior work showing that the ability to plan ahead to solve problems develops rapidly during preschool years. Second, we also found that the proportion of children who finished the task with maximal efficiency varied across conditions. This pattern is also reasonable given that the Standing Bridge required a more principled, planned approach for success; due to the constraint of gravity, there was only one viable decomposition solution whereas the Sideways Bridge could be built in a few different ways. Third, we found an age by condition interaction in measures of efficiency (time-to-completion and number of coins used); however, the difference between conditions in efficiency increased with age, rather than decreasing with age. In other words, only the older children showed the expected difference between conditions. This suggests that the task was generally quite difficult for young children; even though 4- and 5-year-olds still successfully passed several practice trials and understood the task instructions, many of them struggled to complete the task in both conditions.

Collectively, these data provide an informative window into how children engage in problem decomposition to solve a complex task. Older children's near-perfect performance in the Sideways Bridge condition suggests that they were able to create substructures and use them to assemble the bridge correctly. Thus, the primary challenge children faced in this task may have been identifying and generating a plan to construct the "correct" components prior to building, especially when there was just one solution (Standing Bridge condition).

The fact that younger children struggled in both conditions raises questions about whether preschool-aged children suffer from a genuine lack of ability to engage in problem decomposition. However, the results do not allow us to directly explore this possibility because the task in Experiment 1 was open-ended and required children to engage in all aspects of problem decomposition—generating, evaluating, and ex-

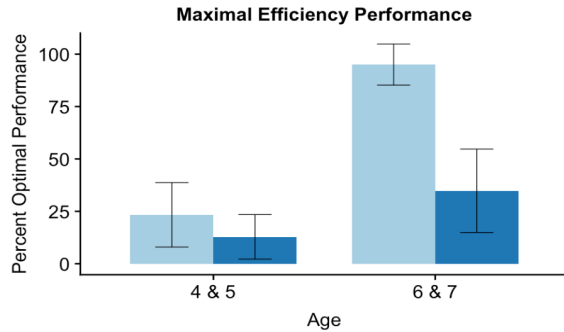


Figure 3: Percentage of children who completed the task with maximal efficiency (only used 3 coins), collapsed into age bins. Error bars are bootstrapped 95% CI.

ecuting solutions. However, there are reasons to believe that, when the demands for generating the plans are removed, even younger children can successfully *evaluate* the viability of a given plan. Compared to the process of generating a plan for decomposing a problem or a structure, evaluating the viability of an existing plan is arguably an easier task. Prior work suggests that preschool-aged children can easily assemble structures (Cortesa et al., 2018) and evaluate the relative difficulty of building different structures (Gweon et al., 2017), suggesting that even though the younger children in Experiment 1 (4- and 5-year-olds) struggled to generate and execute the plans themselves, they may be capable of evaluating the viability of existing decomposition plans.

In Experiment 2, we test this hypothesis with a simple binary-choice paradigm where we asked children to choose one of two pre-generated plans (i.e., choose the plan that would result in a self-supporting structure). Given the simplicity of the task, in addition to 4- and 5-year-olds, who we expected would succeed, we also tested 3-year-olds; while we did not have strong a priori predictions regarding the 3-year-olds' performance, having a broader age group would allow us to capture the developmental trajectory of this ability.

Experiment 2

Methods

Participants A total of 78 children were recruited from a local children's museum and a university-affiliated preschool (28 3-year-olds: $M = 3.51$ (3.02–3.98); 26 4-year-olds: $M = 4.42$ (4.01–4.93); 24 5-year-olds: $M = 5.52$ (5.03–5.93)). An additional 12 children were dropped from analyses because: (1) they did not speak English ($N=1$), (2) they did not complete the study ($N=4$), (3) they failed the pretrial task ($N=3$, see Procedure), (4) parents interfered ($N=1$), or (5) the experimenter made an error ($N=3$).

Stimuli Stimuli were similar to the blocks structures used in Experiment 1. A 'T' shaped structure (practice trial) and the upright bridge from Exp. 1 (main trial) were used as target structures. For both trials we prepared two sets of blocks, pre-

configured in the shape of the target structure and laid flat on the surface. Critically, only one of the two sets would result in the correct self-supporting structure (see Fig. 4).

Procedure Children were introduced to block pieces of various shapes. In the practice trial, the experimenter presented the T-shaped block structure along with two potential solutions, and asked: "Can you help me build a new building that looks just like this one and can stand up all by itself? We can use these blocks (pointing to one set) or these blocks (pointing to the other set). Only one will work."

After the child selected one set of blocks, the experimenter allowed the child to use the selected blocks to construct the target structure. Regardless of whether or not the child chose correctly, the experimenter allowed the child to attempt construction with the other set. After the child succeeded in constructing the structure with the correct block set and failed with the incorrect set, the experimenter reiterated that the child could build the building with one set of blocks but not with the other, as it would fall over, so only one set would work. We excluded children who failed to select one of the block options in this pretrial task or who began playing with the blocks before listening to the full explanation.

In the main task, the children were asked to choose one of two solutions that would result in a standing bridge. We marked a child as having made a selection when they physically picked up one of the sets of blocks. The position of the correct solution (L/R) was counterbalanced across subjects.

Results and Discussion

We first ran a logistic regression on children's choice with age as a continuous variable. The effect of age was trending towards significance ($z = 1.81$, $p = 0.07$). We then looked at each age group separately. Three-year-olds' responses did not differ from chance ($M = 0.57$, $CI = 0.39-0.75$), whereas four-year-olds ($M = 0.81$, $CI = 0.65-0.92$) and five-year-olds ($M = 0.83$, $CI = 0.67-0.96$) showed robust success.

To succeed at this task, a child had to be able to understand the constraints applied to the problem (gravity) and physically simulate the stability of the resulting structure to choose the appropriate solution. The success of four- and five-year-old children on this task provides suggestive evidence that they are already capable of such sophisticated physical reasoning. The results also suggest that, even though children in this age group struggled to complete the task in Experiment 1, their difficulty with that task did not stem from an inability to assess the viability of a given plan.




Figure 4: Schematic of two potential solutions presented to children in Experiment 2.




Figure 5: Success rate in Experiment 2. Error bars are bootstrapped 95% CI. Dotted line indicates chance-level.

General Discussion

Our goal was to assess whether the basic capacities for problem decomposition—one of the key components of CT—are present even early in childhood. In Experiment 1, we used a block-building task that involves generating, evaluating, and executing an appropriate decomposition plan to build a physical structure. The results suggest that capacities for top-down design and problem decomposition continue to develop well past the preschool years, and that children become more successful and efficient with age. Although many younger participants failed to complete the task in Experiment 1, in Experiment 2 we find evidence for one of the key steps in successful problem decomposition: children as young as age 4 were able to evaluate the viability of potential solutions.

Experiment 1 featured a rather complex task with high verbal demand for understanding the instructions, which may have increased the task load. Furthermore, this study jointly required top-down design to generate appropriate solutions, evaluation of those solutions, and the actual execution of the plan to assemble the components. Children’s struggle with this task could reflect their difficulties in any or all of these steps. Experiment 2 isolates one particular aspect of problem decomposition. Results suggest that 4- and 5-year-old children can compare and evaluate two different decomposition solutions and select the correct one. These results complement prior work (Cortesa et al., 2018) which showed that young children can construct target structures from predetermined components; beyond using a given set of components to assemble the target structure, our results show that 4-year-olds are able to reason ahead under the constraints of a task to infer the correct set of components, even before they engage in actual assembly. Collectively, these findings indicate that some basic underlying capacities for problem decomposition may begin to emerge in preschool years, but they also continue to develop well beyond this age.

One might wonder whether children’s abilities to engage in problem decomposition in our task is restricted by the physical/spatial domain. Prior work indicates that the ability to engage in basic spatial reasoning emerges early in life (Newcombe & Huttenlocher, 2003). For instance, 5-year-olds show successful mental rotation of a paper cut-out object on

a 2-D plane (Frick, Hansen, & Newcombe, 2013) and understand how a scene would look from another person’s perspective (Borke, 1975). Our results suggest that even 4-year-olds can mentally rotate a 3-D structure to assess its stability.

Our study focused on a concrete problem with a clear visual representation. Our tasks were intentionally reflective of a thinking pattern common to programming. To solve a programming problem a programmer must identify independently solvable pieces, construct them separately, and then recombine them into a cohesive solution. Of course, throughout this process, the programmer must weigh constraints to make decisions about optimal components or solutions. Similarly, in Experiment 1, children had to identify, construct, and reassemble components of a larger physical structure; there was no possible way to build it directly. Thus, one important question is whether the ability to decompose a problem in the spatial domain extends to more abstract CT problems. Future work might ask whether children’s success in this task transfers to decomposition of larger tasks in other STEM areas, such as programming. One possibility is that training children to engage in decomposing a physical structure might also help them decompose a larger programming problem.

Another interesting avenue for future exploration is that the use of concrete objects in physical space might make it easier for children to engage in successful decomposition even in these more abstract domains. Indeed, adults often transform complex abstract tasks into concrete forms, such as diagrams, to avoid trial-and-error in a complex project. We look forward to future work that asks whether physical affordances and manipulatives support children’s abstract problem solving in a similar way.

Mark Guzdzial, a leading researcher in CS education, wrote: “An open research question is what an elementary school child can learn about computing and what should be taught at what ages” (2015). Our work, along with prior research in cognitive development, suggests that CT is not a unitary construct that emerges at any single age. It involves a range of mental operations which may involve independent developmental trajectories. While children might be able to identify flaws in systems or construct those systems from predetermined parts in preschool, they may not develop the ability to generate those parts until much later. Capacities underlying other CT concepts, such as abstraction, data representation, or parallelization, likely also develop in a piecemeal manner that remains to be discovered.

We look forward to more future work that bridges the gap between cognitive development and CS education research. Our work here represents a first step at demonstrating children’s developing capabilities in a critical component of computational thinking: problem decomposition. We show that children may be able to learn basic computational thinking skills as early as preschool, but that these capacities continue to develop well into elementary years. As educators continue to develop CS curriculum, these results can inform when and how to teach early programming concepts to young students.

Acknowledgements

Thanks to Rhonda Sandifer for her help with data collection. We also thank staff, children, and families at Bing Nursery School and at the Palo Alto Junior Museum and Zoo.

References

- Barr, V., & Stephenson, C. (2011). Bringing computational thinking to k-12: what is involved and what is the role of the computer science education community? *ACM Inroads*, 2(1), 48–54.
- Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, 201306572.
- Borke, H. (1975). Piaget's mountains revisited: Changes in the egocentric landscape. *Developmental Psychology*, 11(2), 240.
- Brennan, K., & Resnick, M. (2012). New frameworks for studying and assessing the development of computational thinking. In *Proceedings of the 2012 annual meeting of the american educational research association* (Vol. 1, p. 25).
- Brosnan, M. J. (1998). Spatial ability in children's play with lego blocks. *Perceptual and Motor Skills*, 87(1), 19–28.
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers exploratory play. *Cognition*, 120(3), 341–349.
- Cooper, S., Wang, K., Israni, M., & Sorby, S. (2015). Spatial skills training in introductory computing. In *Proceedings of the eleventh annual international conference on international computing education research* (pp. 13–20).
- Cortesa, C., Jones, J., Hager, G., Khudanpur, S., Landau, B., & Shelton, A. (2018). Constraints and development in childrens block construction.
- Frick, A., Hansen, M. A., & Newcombe, N. S. (2013). Development of mental rotation in 3-to 5-year-old children. *Cognitive Development*, 28(4), 386–399.
- Gardner, W., & Rogoff, B. (1990). Children's deliberateness of planning according to task circumstances. *Developmental Psychology*, 26(3), 480.
- Gergely, G., Nádasdy, Z., Csibra, G., & Bíró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, 56(2), 165–193.
- Gopnik, A. (2012). Scientific thinking in young children: Theoretical advances, empirical research, and policy implications. *Science*, 337(6102), 1623–1627.
- Guzdial, M. (2015). Learner-centered design of computing education: Research on computing for everyone. *Synthesis Lectures on Human-Centered Informatics*, 8(6), 1–165.
- Gweon, H., Asaba, M., & Bennett-Pierre, G. (2017). Reverse-engineering the process: Adults and preschoolers ability to infer the difficulty of novel tasks. In *Proceedings of the 39th annual conference of the cognitive science society*.
- Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends in cognitive sciences*, 20(8), 589–604.
- Jara-Ettinger, J., Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2015). Childrens understanding of the costs and rewards underlying rational action. *Cognition*, 140, 14–23.
- Kamps, F. S., Julian, J. B., Battaglia, P., Landau, B., Kanwisher, N., & Dilks, D. D. (2017). Dissociating intuitive physics from intuitive psychology: Evidence from williams syndrome. *Cognition*, 168, 146–153.
- Klahr, D., & Robinson, M. (1981). Formal assessment of problem-solving and planning processes in preschool children. *Cognitive Psychology*, 13(1), 113–148.
- Kramer, J. (2007). Is abstraction the key to computing? *Communications of the ACM*, 50(4), 36–42.
- Legare, C. H. (2012). Exploring explanation: Explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child development*, 83(1), 173–185.
- Liu, S., Ullman, T. D., Tenenbaum, J. B., & Spelke, E. S. (2017). Ten-month-old infants infer the value of goals from the costs of actions. *Science*, 358(6366), 1038–1041.
- Newcombe, N. S., & Huttenlocher, J. (2003). *Making space: The development of spatial representation and reasoning*. MIT Press.
- Schulz, L. (2012). The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in cognitive sciences*, 16(7), 382–389.
- Scott, R. M., & Baillargeon, R. (2013). Do infants really expect agents to act efficiently? a critical test of the rationality principle. *Psychological science*, 24(4), 466–474.
- Verdine, B. N., Golinkoff, R. M., Hirsh-Pasek, K., Newcombe, N. S., Filipowicz, A. T., & Chang, A. (2014). Deconstructing building blocks: Preschoolers' spatial assembly performance relates to early mathematical skills. *Child development*, 85(3), 1062–1076.
- Wai, J., Lubinski, D., & Benbow, C. P. (2009). Spatial ability for stem domains: Aligning over 50 years of cumulative psychological knowledge solidifies its importance. *Journal of Educational Psychology*, 101(4), 817.
- Wing, J. M. (2006). Computational thinking. *Communications of the ACM*, 49(3), 33–35.
- Yildirim, I., Gerstenberg, T., Saeed, B., Toussaint, M., & Tenenbaum, J. (2017). Physical problem solving: Joint planning with symbolic, geometric, and dynamic constraints. *arXiv preprint arXiv:1707.08212*.
- Zhang, M., Igarashi, Y., Kanamori, Y., & Mitani, J. (2017). Component-based building instructions for block assembly. *Computer-Aided Design and Applications*, 14(3), 293–300.