

Technische Universität München, School of Education

Active learning strategies in child- and adulthood: an interdisciplinary perspective

Angela Jones

Vollständiger Abdruck der von der TUM School of Education
der Technischen Universität München zur Erlangung des akademischen Grades
eines Doktors der Philosophie (Dr. phil)
genehmigten Dissertation.

Vorsitzende/-r: Prof. Dr. Christina Seidel

Prüfende/-r der Dissertation:

1. Prof. Dr. Azzurra Ruggeri

2. Prof. Dr. Anna-Julietta Baumert

Die Dissertation wurde am 22.03.2021 bei der Technischen Universität München
eingereicht und durch die TUM School of Education am 22.06.2021 angenommen.

Contents

1	Acknowledgements	3
2	Abstract	4
2.1	English	4
2.2	Deutsch	5
3	Introduction	7
3.1	The many faces of active learning: different definitions of active learning influence research methods and findings	8
3.2	The development of active learning strategies	11
3.3	Active versus passive learning	16
3.4	Applied active learning: education and scientific reasoning	19
3.5	This dissertation: research goals	21
3.6	Other potential research directions	22
4	General Methodology	24
4.1	Samples	24
4.2	Interdisciplinary approach	24
5	Summary of associated published manuscripts	27
5.1	Paper 1: How is the hypothesis space represented? Evidence from young children’s active search and predictions in a multiple-cue inference task	28
5.2	Paper 2: Active function learning	29
6	General Discussion	32
6.1	Future directions	36
6.2	Conclusion	39

7	References	40
8	Appendix	52

1 Acknowledgements

Writing a dissertation is a difficult yet exciting journey, during which one grows both as a person and a scientist, surrounded, influenced and supported by many inspiring people. During my journey, I not only developed my research skills but was also able to reach, and extend, the limits of my capabilities. I would like to thank all the people who accompanied me on this journey.

First, I thank my advisor, Prof. Dr. Azzurra Ruggeri, for her guidance and support, and for providing me with the opportunity to enter the world of research and a warm, bustling working environment. She was always available when I needed input but also gave me room to grow and take initiative.

Second, I also thank my mentor and co-author, Dr. Eric Schulz, for his advice and patient tutoring in the finer points of computational modeling. I would also like to thank everyone who helped me implement my studies: Prof. Dr. Doug Markant, Dr. Thorsten Pachur and Dr. Neil Bramley, for their computational expertise and helpful input, Federico Meini, for his invaluable programming skills which allowed me to implement my tasks as game apps and record very clean data, and Teresa Niemann, Silvia Martín Lence, Katharina Gräfin zu Lynar, and Nele Metzner, who did the testing for my studies and were a pleasure to work with at all times.

I am also grateful to my friends and fellow PhD students, Nora Swaboda and Costanza De Simone, without whom my time at the institute would undoubtedly have been much less fun.

Finally, I thank my 8th grade English literature teacher, Mrs. Slaiding, for telling me to give up on a scientific career path because my mathematics grades were not high enough. Ignoring her well-meaning advice was one of the best decisions I have ever made and has taught me that while advice is valuable, it is not as valuable as my own perseverance and hard work in pursuing my goals.

2 Abstract

2.1 English

Active learning has been extensively studied in both adults and children, using a wide variety of tasks. Studies focusing primarily on how children's learning strategies change have shown substantial developmental improvements in search efficiency, particularly between the ages of 5 and 10. However, studies focusing on children mostly employ behavioral paradigms, which provide little insight into the processes underlying children's search strategies and giving rise to these improvements. In adults, computational modeling has proven to be a successful approach for addressing this question, showing that adults' search is guided by the learners' hypothesis space, and that their strategies aim to systematically reduce uncertainty about which hypotheses are correct. Several studies have also demonstrated learning advantages from active over passive learning, but others contradict these results. This dissertation project aims to further our understanding of these processes by combining computational and behavioral methods to study the active learning strategies employed by 5- to -7-year-old children, and adults. The results of this dissertation show for the first time that the hypothesis space can change during a task, with developmental differences in its structure emerging between the ages of 5 and 7, and that sampling strategies can be transferred between tasks to a certain extent, a fact which is unaffected by learning condition. These studies provide important insights into how computational processes affect the implementation of active learning strategies in different tasks, and how developmental changes in these processes may impact children's strategies. Together, these studies provide a deeper understanding of how computational and cognitive processes may interact to give rise to documented developmental changes in strategy use, and help clarify existing debates about the benefits of active learning. Insights from this line of research also have potential for building up a scientific framework to guide instructors and educational app creators in designing

interventions.

2.2 Deutsch

Strategien des aktiven Lernens sind sowohl bei Erwachsenen als auch bei Kindern unter Verwendung einer Vielzahl von Aufgaben eingehend untersucht worden. Studien, die sich mit der Veränderung der Lernstrategien von Kindern befassen, haben gezeigt, dass sich die Sucheffizienz besonders im Alter von 5 bis 10 Jahren erheblich weiterentwickelt und verbessert. Studien mit Kindern verwenden jedoch meist Verhaltensparadigmen, die nur wenig Einblicke in die Prozesse geben, die den Suchstrategien von Kindern zugrunde liegen oder zu entwicklungsbedingten Verbesserungen beitragen. Bei Erwachsenen hat sich die Computermodellierung als erfolgreicher Ansatz zur Beantwortung dieser Fragen erwiesen. Ergebnisse aus solchen Untersuchungen zeigen, dass die Informationssuche vom Hypothesenraum der Lernenden geleitet wird und dass ihre Strategien darauf abzielen, die Unsicherheit darüber, welche Hypothesen korrekt sind, systematisch zu verringern. Mehrere Studien haben auch Lernvorteile von aktivem gegenüber passivem Lernen gezeigt, wobei jedoch auch widersprüchliche Daten existieren. Ziel dieses Dissertationsprojekts ist es, den kindlichen Suchstrategien zugrundeliegenden Prozessen und der Ursprünge entwicklungsbedingter Veränderungen besser zu verstehen. Dafür werden die aktiven Lernstrategien von 5- bis 7-jährigen Kindern und Erwachsenen mit einer Kombination aus computergestützten und verhaltensbasierten Methoden untersucht. Diese Dissertation zeigt, zum ersten Mal, dass sich der Hypothesenraum während einer Aufgabe ändern kann, und dass seine Struktur zwischen den Ältern von 5 und 7 entwickelt. Ich zeige auch, dass aktive Stichprobenstrategien zwischen verschiedenen Arten von Aufgaben übertragen werden können. Diese Studien liefern wichtige Erkenntnisse darüber, wie kognitive Prozesse die Implementierung aktiver Lernstrategien in verschiedenen Aufgaben beeinflussen und wie sich entwicklungsbedingte Veränderungen in diesen Prozessen auf kindliche Suchstrategien auswirken.

Zusammen tragen die Ergebnisse dieser Dissertation dazu bei, ein tieferes Verständnis dafür zu erlangen wie rechnerische und kognitive Prozesse zusammenwirken und dokumentierte entwicklungsbedingte Veränderungen in der Strategieverwendung bedingen. Die Erkenntnisse aus dieser Forschungsarbeit können dabei helfen, bestehende Debatten über die Vorteile aktiven Lernens zu informieren und besitzen Potential für die Erarbeitung eines wissenschaftlichen Rahmenwerks, welche Pädagogen sowie Entwickler von Bildungs-Apps in der Gestaltung neuer Bildungsangebote leiten und unterstützen könnte.

3 Introduction

Active learning has been an enduring buzzword in both the mainstream media and scientific communities for some time. Entering the keywords ‘active learning’ into Google generates about 2.530.000.000 search results, and the same query entered into Google Scholar returns approximately 5.480.000 scientific publications (approximately 7.070.000 results for the similar term ‘information search’), reflecting the fact that the scientific community has been investigating active learning for over 50 years, and in a huge variety of fields, from machine learning, to cognitive science, to educational science. This is therefore an important research topic with extremely varied applications, the relevance of which continues to increase with educators’ and parents’ growing interest in optimizing the classroom learning experience.

Active learning is widely considered to be one of the best approaches to learning in education. Several educational theories advocate a different idea of what it means to be an active learner and the benefits of such learning relative to passive instructional methods, such as discovery learning (Bruner, 1961), experiential learning (Kolb, 2014), inquiry learning (Kuhn, Black, Keselman, & Kaplan, 2000), constructivism (Steffe & Gale, 1995), and self-regulated learning (Boekaerts, 1997). These theories vary in their focus but share the belief that active learning leads to improved learning outcomes. This claim is often supported in comparisons with more traditional forms of passive learning such as lecture-based teaching (Bonwell & Eison, 1991; Freeman et al., 2014). However, despite the popularity of active learning, it is not always clear why it works, or does not, in real world settings. This is because active instruction usually differs from passive learning in several respects, and it is often unclear which of these differences lead to observed results. In fact, the concept of active learning now encompasses a large variety of instructional techniques, which usually refer to a combination of physical activity or interaction, the creation or explanation of learning materials, planning learning activities, question asking, metacognition, and social collaboration. Furthermore,

children’s active learning strategies have not always been characterized in detail, and the relationship between strategy-specific factors such as sampling strategies (i.e., how learners choose what information to target at each step of the search) and cognitive factors such as executive functions remains unclear. Together, the variable definitions of active learning and the knowledge gaps regarding its development make it difficult to identify what causes differences in performance during active learning, both across development and between learning conditions, and to predict whether these effects can be generalized to other kinds of activities or materials.

The present dissertation focuses on active learning in the domain of cognitive science. Its aim is to deconstruct children and adults’ active learning strategies in order to reach a deeper understanding of how these strategies are implemented in child- and adulthood, and identify which factors potentially drive developmental changes. To this end, I combine computational modeling with behavioral paradigms to study how two crucial aspects of active learning strategies — the hypothesis space and active sampling strategies — are implemented at the computational level in 5- to 7-year-old children and adults. Note that active learning will be defined as self-led or self-regulated learning (wherein the learner controls what to learn, and when) in the present work, and used interchangeably with the term ‘information search.’

3.1 The many faces of active learning: different definitions of active learning influence research methods and findings

It is not only the ongoing debates over the merits of active learning versus guided instruction (Kirschner, Sweller, & Clark, 2006; Mayer, 2004; Prince, 2004) which add to the difficulty of understanding which factors drive developmental changes in active learning strategies, but also attempts to create precise taxonomies of active behaviors and their predicted effects on learning (Chi, 2009; Chi & Wylie, 2014). For instance, Chi (2009) proposed several gradations of what is commonly considered “active” learn-

ing: active, constructive, and interactive. She defines *active* learning as the lowest level, signifying merely doing something, for example highlighting and underlining relevant sentences in written material, while *constructive* learning is a step above and purported to require a secondary output such as summarizing material, drawing concept maps, and self-explaining. Finally, *interactive* learning is the most engaged form of learning and leads to the best learning outcomes, as it involves talking to another person, interacting with a system like a virtual tutoring system, and physical interactions. This framework is best suited to an educational context as Chi's definitions of active behavior all involve actions that are likely to take place in a classroom. This also highlights an additional source of confusion when it comes to defining active learning: the context of the discussion. As Chi notes in her 2009 work, behavior such as repeating a set of words to be memorized would be categorized as active in her framework, but passive in the memory literature (p. 76).

The debate over whether instruction or active learning is better for learning outcomes is also complicated by the varying definitions of "active learning". A good example of this is in Klahr and Nigam (2004). Students were either asked to "discover" on their own, given a set of materials, the Control of Variables Strategy (CVS), an approach to causal learning which involves isolating the effects of a single variable at a time on a causal system (Kuhn & Brannock, 1977, discovery learning, or active condition), or observed as an experimenter conducted different experiments using this strategy and explained why this was an effective approach (direct instruction, or passive condition). The authors found that students in the direct instruction condition outperformed those in the discovery learning condition, suggesting that in this case, passive learning was more effective than active learning. However, the instructions given to students in the direct instruction condition included a prompt to think about how sure they could be that a variable had an effect on an outcome as a result of each experiment. According to Chi's (2009) framework, this would therefore be considered a constructive learning

task, not a passive one. Therefore, as indicated by Chi (2009), Klahr and Nigam's (2004) finding that students in this condition displayed better learning outcomes than those in the discovery learning condition is also consistent with her prediction that constructive learning leads to better learning outcomes than active learning. This example illustrates that fact that studies purporting to contrast passive and active learning may not necessarily be interpreted the same way, depending on the framework used.

Outside of the educational literature, active learning has been studied in a wide variety of tasks which tend to capture skills which are important in ubiquitous, real-world learning tasks, in both adults and children. Here, the focus has been more on learners' search patterns themselves, with learning outcomes being a secondary concern, and little consideration of the transfer of strategies between tasks. However, there are some methodological differences between studies adopting a developmental perspective and those focused primarily on adults. Studies investigating active learning in children employ mainly behavioral methods, although there has been growing interest in combining these methods with computational modeling (e.g., Mata, von Helversen, & Rieskamp, 2011; Ruggeri, Sim, & Xu, 2017; von Helversen, Mata, & Olsson, 2010). Studies of adult learning have, in contrast, been successfully using computational methods for some time (e.g., Enkvist, Newell, Juslin, & Olsson, 2006; Juslin, Jones, Olsson, & Winman, 2003; Juslin, Olsson, & Olsson, 2003; Markant & Gureckis, 2014; Markant, Settles, & Gureckis, 2016; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003), leading to the elaboration and refinement of computational-level concepts such as the hypothesis space (i.e., the mental structure of representations of task-relevant information) and sampling strategies (i.e., what learners care about when selecting which information to look at during their search).

For instance, Markant and Gureckis (2014) modeled adults' active learning strategies in a category learning task. Participants were presented with "loop antennas" for televisions, differing in size and the angle of a central diameter, and had to learn which

of two fictional channels each antenna received based on these two cues, and could change the values of these dimensions to see which channel was received as a result of the changes. Examination of participants' information search strategies showed that they tended to focus their search on training items that were close to the category boundaries, where there was more uncertainty about how to classify each item, and preferential use of this strategy was associated with higher accuracy. Furthermore, adults tend to break learning problems into smaller "chunks", and use this approach to resolve uncertainty between two hypotheses at a time in each chunk (Markant et al., 2016). Quantifying the efficiency of adults' strategies has shown that their choices are informative and generally aim to reduce uncertainty at every step (i.e., they do not sample randomly but choose to look at information based on how useful it is for differentiating between competing hypotheses; e.g., Gureckis & Markant, 2009; Steyvers et al., 2003), an effect which is amplified by restricting how much information they can look at during a learning phase (Steyvers et al., 2003). These findings illustrate two important characteristics of information search strategies: that they are guided by the hypotheses entertained by each individual learner (i.e., the hypothesis space), and that they often aim to reduce uncertainty about the solution to a learning problem. These examples also illustrate how computational modeling of active learning strategies can provide important insights into some of the processes underlying human learning behavior.

3.2 The development of active learning strategies

In line with these findings, a growing body of literature demonstrates that even infants prefer to explore more uncertain options (L. Schulz, 2015; Stahl & Feigenson, 2015), and that this preference is present throughout development, with preschoolers also being more likely to explore when presented with confounded evidence, i.e., when they are uncertain about the causal mechanism at work (Cook, Goodman, & Schulz, 2011; L. E. Schulz & Bonawitz, 2007), and when they witness evidence which violates their

beliefs (Bonawitz, van Schijndel, Friel, & Schulz, 2012; Legare, Gelman, & Wellman, 2010). Moreover, using a task in which noisy rewards were spatially correlated on a grid, E. Schulz, Wu, Ruggeri, and Meder (2019) showed that 7- to 11-year-old children learned more slowly than adults and generalized less, but explored more (although not randomly) and that they preferentially explored areas with high uncertainty, suggesting that they were even more strongly motivated than adults to reduce uncertainty.

The idea that children preferentially explore under conditions of uncertainty is further supported by work on curiosity (e.g., see Jirout & Klahr, 2012; Kidd & Hayden, 2015; Kidd, Piantadosi, & Aslin, 2012). For instance, Information Gap Theory (Loewenstein, 1994) proposes that curiosity is internally motivated, in that it arises when an individual becomes aware of a gap in their knowledge, i.e., when they are uncertain about something. Awareness of this knowledge gap induces a desire to reduce it, which is resolved by looking for the missing information. Jirout and Klahr (2012) further refined this idea by proposing that curiosity is a “threshold of desired environmental uncertainty that leads to exploratory behavior” (p. 127), which is consistent with evidence that human beings generally seek to resolve uncertainty, but does not account for an upper bound of “desired uncertainty”. Even in infants, curiosity and attention are focused on stimuli or information that are moderately unfamiliar, rather than completely familiar or completely unfamiliar, leading to an inverted-U-shaped curve (e.g., Kang et al., 2009; Kinney & Kagan, 1976). Kidd et al. (2012) termed this the “Goldilocks Effect”, in reference to the fact that there is a level of uncertainty that is “just right”—neither too little or too much — which elicits curiosity, and therefore exploration. However, it is worth noting more recent work by van Schijndel, Jansen, and Raijmakers (2018) which has shown that, at least in the context of scientific causal learning, curiosity is not related to the quality of information search strategies, but rather to knowledge acquisition. Therefore, the relationship between curiosity and active learning is not as straightforward as it might first appear. Interestingly, the authors posited that these re-

sults suggested the existence of two parallel processes for inquiry-based learning: one which deals with the planning of experiments, and the other with the later evaluation of and reflection on those experiments. If true, this may complicate the interpretation of results linking active learning strategies to learning outcomes. On the other hand, this study focused on the kinds of causal learning that take place in the science classroom, which are a special case of active learning, so it remains to be seen whether these results might extend to other learning contexts.

However, despite this early sensitivity to uncertainty, studies specifically investigating how *effective* children's active learning strategies are have shown that young children search inefficiently; their search strategies develop from seemingly undirected and exhaustive exploration to more adult-like levels of efficiency, with particularly large improvements observed between the ages of 5 and 10 (e.g., Davidson, 1991a, 1991b, 1996; Mosher & Hornsby, 1966; Ruggeri, Lombrozo, Griffiths, & Xu, 2016; Ruggeri et al., 2017). In question-asking tasks, participants' task is to find the correct answer from a series of options, or hypotheses, using as few yes-or-no questions as possible. Thus, two types of questions are possible: hypothesis-scanning, which target one option at a time (e.g., is it the daisy?), and constraint-seeking, which target several options at once (e.g., is it a plant?). When all options are equally likely, using constraint-seeking questions is the most efficient approach, because it allows the question asker to narrow down the hypotheses quickly to find the correct one, and this is adults' default approach. In contrast, children under the age of 7 almost exclusively rely on hypothesis-scanning questions, and only begin to transition to constraint-seeking questions between the ages of 7 and 10 (Herwig, 1982; Mosher & Hornsby, 1966; Ruggeri & Feufel, 2015; Ruggeri et al., 2016).

Likewise, in information board procedures, where information about different alternatives is uncovered sequentially before choosing one of the alternatives, younger children (6- and 7-year-olds) search more of the available information than older chil-

dren (9- to 13-year-olds; Betsch, Lang, Lehmann, & Axmann, 2014; Davidson, 1991a, 1991b, 1996) and seem to have difficulty identifying which pieces of information are more relevant or important for the decision (e.g., Betsch et al., 2014; Davidson, 1996, but see von Helversen et al., 2010). However, the factors which drive these changes, as well as the hypotheses that children use to guide their search at different ages and exactly how they take uncertainty into account, remain unknown. Taking inspiration from the adult literature and adopting computational methods to address these questions would help deepen our understanding of documented changes in learning behaviors and shed light on these questions.

In order to be truly efficient, active learning strategies must also be flexible, tailored to one's learning environment. Different strategies vary in informativeness depending on the characteristics of the task at hand, as well as on the previous knowledge and expectations of the searcher (Todd & Gigerenzer, 2012). Being able to adapt one's learning strategies to the current learning context, an ability referred to as ecological learning (Ruggeri & Lombrozo, 2015; Ruggeri et al., 2017), is therefore crucial for maximizing learning effectiveness. Adults are ecological learners, and adapt their learning strategies according to sparsity (i.e., the number of hypotheses affecting an outcome) in a number of non-causal hypothesis testing tasks (Hendrickson, Navarro, & Perfors, 2016; Langsford, Hendrickson, Perfors, & Navarro, 2014; McKenzie, Chase, Todd, & Gigerenzer, 2012; Navarro & Perfors, 2011; Oaksford & Chater, 1994). For example, Hendrickson et al. (2016) showed that people switched from requesting positive to negative examples of a concept when the overall proportion of positive cases increased. This also holds true for causal learning tasks, as Coenen et al. (2019) showed that adults also adapt their learning strategies to causal sparsity.

However, perhaps surprisingly, adults do not seem to be better ecological learners than children. For example, in question-asking games, 7-year-olds are equally able to tailor their questions to the statistical structure of the environment (Ruggeri & Lom-

brozo, 2015), and even seem to do so more readily than adults when the most efficient question-asking strategy is not the one that adults employ by default (Ruggeri & Lombrozo, 2015). This is consistent with evidence that sensitivity to environmental probabilities emerges very early on in life. Infants as young as 10 to 12 months old are not only sensitive to probabilities, but also use probabilistic information to make judgements and predictions and revise them after observing new evidence (Denison & Xu, 2014; Gweon, Tenenbaum, & Schulz, 2010; Kushnir & Gopnik, 2005). In fact, assessments of probabilistic cognition in two Mayan indigenous groups with no formal education suggests that this form of cognition may be a fundamental, and universal, human skill, one which does not require any particular training to acquire (Fontanari, Gonzalez, Vallortigara, & Girotto, 2014). From the age of 5 or 6, children have also been shown to integrate prior probabilities with feedback and subsequent evidence (Denison, Reed, & Xu, 2013; Girotto & Gonzalez, 2008; Gonzalez & Girotto, 2011, but note that Gonzalez and Girotto found that children need additional instruction to properly refine this skill) and make inferences that are consistent with the general principles of Bayesian inference (e.g., Eaves & Shafto, 2012; L. E. Schulz, Bonawitz, & Griffiths, 2007). Furthermore, young children have been shown to correctly infer unusual causal relationships faster than adults do, suggesting that they are capable of adapting their learning process to uncommon environmental structures more readily than adults (Gopnik, Griffiths, & Lucas, 2015). In sum, sensitivity to probabilities and the statistical properties of a task is a fundamental building block of ecological learning which appears to be in place from infancy, making it likely that ecological learning is well within young children's capabilities.

Indeed, recent work with 3- to 5-year-olds has shown that children as young as 3 can adapt their exploratory actions to learning environments with different statistical properties as long as the task does not require them to generate informative questions (Ruggeri, Swaboda, Sim, & Gopnik, 2019). This body of work shows that even though

very young children are inefficient learners, they already use probabilistic information about their learning environment to decide whether, how, and how much to explore: they are therefore not only active, but ecological learners.

3.3 Active versus passive learning

Another question which has been of great interest to the scientific community is whether active learning is better than passive learning. Active learning has mostly been shown to benefit performance in category and causal learning tasks (e.g., Gureckis & Markant, 2012; Steyvers et al., 2003), as well as in memorization tasks (Ruggeri, Markant, et al., 2019; Voss, Gonsalves, Federmeier, Tranel, & Cohen, 2011). This advantage could be explained by a hypothesis-dependent sampling bias, whereby each participant considers different hypotheses at each step of the search, but only the active learners can guide their search accordingly, enabling them to learn the category boundaries or underlying causal relationship more successfully (Markant & Gureckis, 2012; Steyvers et al., 2003). Indeed, even such small-scale forms of control as controlling the order and pacing of study material lead to improvements in memory for the studied material (Harman, Humphrey, & Goodale, 1999; Voss et al., 2011). Therefore, being able to control one's learning material can be a more effective learning strategy than simply observing data.

Similar advantages of active over passive learning have also been demonstrated in spatial exploration. In the spatial learning domain, active exploration involves planning where to go next based on a certain goal. This kind of decision making has been shown to be sufficient to enhance learning of the environment, even in the absence of physical interaction or control of movement. For example, Plancher, Barra, Orriols, and Piolino (2013) compared active drivers and yoked passengers in a virtual driving experiment. Active participants were assigned to one of two conditions: an interaction condition, in which they drove a car along a route dictated by the experimenter, and a planning condi-

tion, in which they decided which direction to turn at each intersection and their choices were carried out by the experimenter. Compared to passive observation, in which participants simply watched a video of the driving experience generated by participants in the interaction condition, both active conditions led to better memory for the layout of the virtual environment and the route taken. Moreover, performance in the planning condition was higher than in the interaction condition, suggesting that deciding where to go enhanced memory independently of the physical act of exploring. This advantage of active over passive learning in spatial navigation tasks also exists in children from the age of 3 (Feldman & Acredolo, 1979; McComas, Dulberg, & Latter, 1997; Poag, Cohen, & Weatherford, 1983). However, the opposite pattern was found for recognition memory of objects encountered along the route, with passive observers showing better recognition relative to both active conditions (Plancher et al., 2013). This suggests that the memory benefits of active exploration are specific to information that is relevant to making exploratory decisions, whereas incidental memory for goal-irrelevant information could be impaired or unchanged relative to passive observation.

Further research is needed to investigate whether this is also true of active learning problems that do not involve spatial exploration. Indeed, while these findings raise the possibility that memory benefits associated with active learning could be restricted to information which falls strictly under the scope of the learner's attention, it may also be possible that some kinds of learning strategies which apply to different contexts could trigger deeper processing of *all* the information available to the learner. For instance, in question-asking tasks, asking constraint-seeking questions requires abstracting *all* the option categories, in order to isolate one group of objects to ask about. In contrast, hypothesis-scanning questions only target one option at a time, and therefore do not necessarily require any deeper processing. It may be that such deeper processing can enhance memory for all the options encountered in this kind of task, rather than only the ones that were explicitly asked about. In this scenario, constraint-seeking questions

would lead to improved memory for the objects encountered in question-asking tasks than hypothesis-scanning questions, unless one was unlucky and had to ask about every single option before finding out that the last one remaining was the correct answer. However, the relationship between different kinds of active learning strategies and learning outcomes has yet to be explicitly investigated.

Many studies find support for the intuition that active learning is better than passive learning, but this is not always the case. For instance, active learners performed better than passive learners in some multiple-cue inference tasks, but worse in categorization tasks (Enkvist et al., 2006; Juslin, Jones, et al., 2003; Juslin, Olsson, & Olsson, 2003). This was explained by active participants in the categorization tasks implementing a learning process more suited for multiple-cue judgement tasks, which did not capture the underlying task structure of the categorization task (Enkvist et al., 2006). Indeed, more recent work has suggested that active learning in categorization and multiple-cue judgement tasks promotes the use of this kind of learning process, called cue abstraction, even when this is not the best approach, whereas observational, or passive, learning does not (Henriksson & Enkvist, 2018). These results suggest that the benefits of active learning may also depend on the specific task, with some kinds of tasks promoting active learning strategies that are less effective than passive observation. Since this work was all carried out with adults, it also raises the question of whether adults may lose some sensitivity to the task structure, perhaps because their prior expectations are strongly ingrained, enough to occasionally override the sensitivity that is present from a young age.

Furthermore, Ruggeri and colleagues (2019) found that the benefits of active learning for memory retention only emerged from age six, and continued to increase until they reached adult-like levels (i.e., an improvement of approximately 5-10% over passive learning) around age ten. This indicates that the relative benefits of active over passive learning can vary across age as well as tasks, and, together with the literature

examining the development of active learning strategies, points to the age range of 5 to 10 years as a period of interest when studying how active learning strategies develop. Further examination of the limits of active learning using a broader range of tasks and a developmental perspective would help clarify under which conditions active learning can be truly beneficial.

3.4 Applied active learning: education and scientific reasoning

Research in education has also produced mixed evidence for the relative benefits of active over passive learning. Comparing active and passive learning is a common approach in studies evaluating educational interventions, especially in the context of scientific reasoning. An important focus of science education research has been to teach the basic principles of how learners should approach such problems in general. Educators have specifically focused on teaching students the principle of isolating or controlling variables, or CVS (i.e., the idea that variables should be tested individually while holding everything else constant), which is a general strategy for approaching many types of causal learning problems and often results in non-confounded evidence (for a review of the control of variables principle, see Zimmerman, 2007). Adults and adolescents, although more likely than young children to use the strategy spontaneously, still show a tendency to sometimes test multiple features at once instead of testing them individually (Kuhn et al., 1995). In contrast, in more complex tasks (with a vast hypothesis space) adults often choose to test one causal relationship at a time by holding most variables at a constant value, perhaps because of a need to reduce the cognitive load (Bramley, Dayan, Griffiths, & Lagnado, 2016). As discussed briefly on p. 9-10, studies comparing direct instruction with active learning (or, as it is usually called in this literature, discovery or inquiry learning) have tended to yield contradictory or simply inconclusive results, with some researchers consistently finding either no difference or an advantage for direct instruction (e.g., Chase & Klahr, 2017; Kirschner et al., 2006; Matlen &

Klahr, 2013; Strand-Cary & Klahr, 2008), while others disagree (e.g., Dean & Kuhn, 2007; Kuhn & Dean, 2005). Overall, the evidence on this question is still unclear.

In the education literature, teaching children the control of variables strategy has been an important focus (e.g., Chen & Klahr, 1999; Kuhn & Angelev, 1976; Kuhn & Brannock, 1977). In fact, mastery of CVS is considered so important for STEM achievement that it features as one of the assessment criteria in national standards for science education in the United States (e.g., see National Academy of Sciences, 2013, p.52). A common finding from empirical studies is that children require extensive training to acquire this strategy and teaching them to transfer it to novel tasks is an even bigger challenge (e.g., Klahr, Fay, & Dunbar, 1993; Kuhn et al., 1995; Kuhn & Phelps, 1982), despite the fact that children already display some precursor skills (Sodian, Zaitchik, & Carey, 1991), and even preschoolers can use CVS somewhat successfully if given careful guidance (van der Graaf, Segers, & Verhoeven, 2015). These difficulties may be partly explained by the fact that this strategy is not always the best possible approach. Indeed, when there are few variables affecting an outcome, testing several hypotheses at once can be a more efficient approach than testing them one by one (Coenen et al., 2019). Children may not readily transfer CVS to other learning problems simply because they are aware, on some level, that this is not the only effective strategy in their toolbox. In sum, studies of active learning in the science education literature have produced interesting but complex and sometimes contradictory results, which are difficult to reconcile without a better understanding of the learning processes involved in different problem-solving approaches, and under different learning conditions. The different possible definitions of active and passive learning within different frameworks and different fields also complicate the interpretation of these studies, as illustrated on p. 9-10.

3.5 This dissertation: research goals

In this chapter, I have provided a high-level overview of the current state of the literature on active learning. Several potential research directions arise from this body of work. First, the active learning strategies employed by children and adults have been characterized to different extents and in various tasks, and information search is known to be guided by the hypothesis space of the learner. However, the exact structure of these representations, as well as the processes and cognitive skills underlying search strategies and driving developmental changes in their use, remain poorly understood. Moreover, it is also unknown to what extent sampling strategies differ between tasks. Second, the differential advantages of active over passive learning are still debated, and a satisfying explanation for these discrepancies has yet to be found due to an incomplete understanding of the cognitive processes that generate the observed behaviours. Third, despite several indications that children are adaptive learners, and may be on par with adults in this respect, the adaptiveness of children's search strategies has seldom been explicitly investigated from a developmental perspective outside of question-asking tasks. As such, it is still unclear to what extent children are adaptive learners and how much this skill develops during childhood.

Thus, the literature reviewed in this chapter has shown that there are three crucial elements of information search strategies that need to be considered when evaluating the implementation of active learning strategies: the hypothesis space, sampling strategies, and ecological learning or adaptiveness. Gaining a better understanding of how these three factors are implemented by children and adults, and how they develop, is very important for identifying potential sources of developmental change and, perhaps, potential targets for interventions that seek to optimize active learning. This dissertation focuses on the former two (the hypothesis space and sampling strategies), for which there is currently the least amount of evidence. In other words, the studies included in this dissertation aim to answer the following main questions:

1. How do learners direct their information search? More specifically, how is their hypothesis space represented at different ages and how closely do their learning strategies map onto the hypothesis space?
2. Do key cognitive skills such as executive functions affect the development of the hypothesis space?
3. What do learners care about when deciding what information to look at next? In other words, how much do their sampling strategies change between different tasks and what specific goals do they try to resolve at each step of the search (e.g., reducing uncertainty or maximising their reward, if the task is remunerated)?

Furthermore, the influence of cognitive skills — which naturally underpin learners’ abilities to plan, execute, and retain information — on the development of all aspects of active learning strategies cannot be ignored. However, this question alone could form the basis of an entire dissertation, and as such is too vast to address here.

3.6 Other potential research directions

For the sake of completeness, this section briefly outlines two other major research directions arising from the current state of the literature on active learning, but which are beyond the scope of this dissertation. First, although children engage in active learning from a very young age, and are known to be prolific question askers and social learners, little attention has been paid so far to how they evaluate *other people’s* active learning skills and use this information to make social judgements. Although there is ample evidence of how children interact with peers and adults to learn new information and identify good informants, it remains poorly understood how children use evidence about how competent other learners are at figuring things out on their own (i.e., engaging in active learning) to infer what other skills or traits they might possess. De Simone and Ruggeri (2020; 2019) found that children begin to selectively generalize an informant’s

competence in searching for information effectively to traits or characteristics that are related to this competence (e.g., being smart or good at solving puzzles) from around age 7, with 7- to 10-year-olds exhibiting adult-like selectivity in these generalizations. In contrast, 3- to 4-year-olds did not draw systematic generalizations, while 5- to 6-year-olds tended to over-generalize the informant's competence to unrelated traits (e.g., being able to see further away). However, beyond this work, the development of this ability is as yet unexplored. In addition, considering the ever-increasing presence of technology accessible to children at home (e.g., robots and virtual assistants like Amazon's Alexa), it is becoming very important to understand more about how children of different ages evaluate these entities as intelligent agents and sources of information. To do so, one must also understand more about how children evaluate the cognitive (and other) traits of other people.

Another avenue of investigation is to empirically evaluate the myriad educational apps and learning robots targeting children (over 80 000 apps are categorized as "for education" on the Apple App store; Apple, 2019), in order to identify training programs that work and those that do not. These tools do not all involve active learning, but many of them use their engaging, interactive and self-paced designs as a selling point and several claim to improve skills such as executive functions, memory and logical reasoning. However, these outcomes have usually not been assessed in peer-reviewed studies (for a review on evidence-based educational apps, see Hirsh-Pasek et al., 2015). Therefore, having a better understanding of the processes underlying children's information search strategies, which training designs are beneficial for enhancing these processes, and how children learn from technological sources could also inform the design of future interventions, allowing them to be properly tailored to the age groups of interest by providing a scientific framework for app creators.

4 General Methodology

4.1 Samples

The child samples in the studies presented here focus on children between the ages of 5 and 7, as this age range is within that identified as a time in which active learning strategies undergo significant transformations. In order to diversify the samples, all children were recruited and tested in various museums in Berlin rather than at schools, after parents gave informed consent (all children were asked for verbal consent as well, and were aware they could withdraw at any time). All children were German or fluent in German (except for one study, where data was collected in the USA so children were all fluent in English; see section 5.1), and did not have any learning disabilities. Adult samples were gathered online, on Amazon Mechanical Turk, and were all fluent in English.

4.2 Interdisciplinary approach

In order to gain deeper insights into learners' search behavior, I combine computational modeling with a variety of behavioral paradigms. Work with adults has already proven the value of this interdisciplinary approach and, as reviewed in the introduction, developmental researchers are also beginning to adopt this method to go beyond purely behavioral data.

The behavioral methods I include in this dissertation are principally non-verbal tasks (in the sense that participants only need enough verbal skills to understand the instructions) which only require limited domain-specific knowledge. These kinds of tasks are well-suited for studying young children's learning because they do not require advanced verbal skills, yet capture learning in ubiquitous real world tasks like multiple-cue inference and causal learning, and are also advantageous because they can be gameified to prevent boredom. The fact that very little domain-specific knowledge was required

to complete them was also important because some learning strategies are known to be affected by the amount of domain-specific knowledge required to implement them properly, particularly in causal learning tasks. For example, implementing CVS correctly in science problems requires specific knowledge of the variables in question and their relationships to each other and the outcome (Edelsbrunner, Schalk, Schumacher, & Stern, 2018). Investigating how the efficiency of learners' strategies changes with the amount of domain-specific knowledge they have is beyond the scope of this dissertation, which focuses on learners' *baseline* active learning skills.

The design of the behavioral tasks was carefully developed to allow for the use of computational modeling on the data. One cannot just combine any behavioral task with computational modeling; a model needs to be tailored to the task as much as the task needs to be structured around the necessity of obtaining data that is useful for a model. The domain of active learning provides particularly fruitful ground for exactly this marriage of disciplines, because such learning tasks usually involve a training phase in which learners make specific learning decisions, the efficiency of which can be quantified, and a performance or prediction phase in which they apply that knowledge, again, usually in the form of specific choices or predictions. This kind of data can comfortably be used by a properly tailored computational model, which can generate predictions of both a learning pattern during the training phase, and subsequent performance decisions in the latter phase. In this dissertation, I present two studies which employ this interdisciplinary approach.

In the first paper, summarized in Chapter 5.1 (Supplement 1; Jones, Markant, Pachur, Gopnik, & Ruggeri, 2021, in press), a multiple-cue inference task is combined with a learning model to investigate how 5- to 7-year-olds' hypothesis space is structured and guides both search and prediction decisions. In this task, children decided which monster pairs to see running in a race, to learn how two cues (color and shape) predicted their relative speed and later bet on the winning monsters. Computational modeling

based on both their learning decisions and their subsequent predictions allowed us to infer from children's behavior how their hypothesis space was structured throughout the entire task and therefore address the first research question. This task also included a memory load manipulation, which allowed us to address the second research question.

The second paper (section 5.2 and Supplement 2; Jones, Schulz, Meder, & Ruggeri, 2018) introduces a novel paradigm in the form of a card game to investigate adults' information search in function learning scenarios and address the third research question. Participants either actively selected or passively observed information to learn about an underlying function connecting two sets of values on the cards, and then had to make predictions about one set of values on some new cards. Here, the computational models were developed to determine what kinds of expectations participants had about the function they had to learn and how this impacted their learning choices, as well as what sampling strategy they used, i.e., what they considered more important when choosing new cards to look at in the active condition, thus addressing my second research question. Building upon the well-established modeling work on human function learning, my co-authors and I developed and compared different variants of rule-based (i.e., models where learning decisions adhere to a strict assumption about the underlying function type, e.g., that it is linear and positive) and non-parametric active learning approaches (i.e., where learning decisions are based on prior assumptions that are less strict than in rule-based models and which are updated according to Bayesian learning) and paired each of them with different plausible sampling strategies to see how well each one captured participants' behavior. Each model was compared both with participants' active learning behavior, and participants' predictions in the prediction phase of the task. This study focused on a linear function as this kind of function is the easiest to learn. However, see Supplement 3 for a follow-up study which further explores how task characteristics such as search horizon (i.e., the number of training choices allowed) and function type influence learning and search.

5 Summary of associated published manuscripts

This chapter summarizes the published manuscripts which cannot be included in their entirety in this dissertation for copyright reasons. Each summary is adapted from the papers' abstracts. There are two such manuscripts: one has been accepted for publication in *Developmental Psychology*, and the other is published in the *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*. Note that this dissertation includes three other manuscripts in the Appendix, two which are currently under review but not yet published, and one published book chapter. These manuscripts provide additional evidence to support the two papers at the core of this dissertation and are referenced when appropriate in the General Discussion (Chapter 6). As the author of this dissertation, I was the first author of all of these publications and therefore had a leading role in the development, implementation, data collection, statistical analysis and writing and submission of these publications.

The first publication was submitted to *Developmental Psychology* in March 2020 and was accepted for publication in February 2021. The full reference is the following:

Jones, A., Markant, D. B., Pachur, T., Gopnik, A., & Ruggeri, A. (2021). How is the hypothesis space represented? Evidence from young children's active search and predictions in a multiple-cue inference task. *Developmental Psychology*.

As the first author, I was responsible for 70% of the work involved in this publication: task elaboration and data collection, statistical analyses, writing the manuscript, responding to reviews, and implementing revisions. Prof. Dr. Douglas B. Markant provided his expertise in computational modeling by building and testing the models presented in the paper (15% of the work). Prof. Dr. Azzurra Ruggeri (10%) and Dr. Thorsten Pachur (10%) took on supervisory roles and provided feedback to guide the development of the paper, as did Prof. Dr. Alison Gopnik (5%).

The second publication was submitted to the conference organizers of the 40th Annual Meeting of the Cognitive Science Society in February 2018 and accepted for pub-

lication in the peer-reviewed conference proceedings in April 2018. The full reference is:

Jones, A., Schulz, E., Meder, B., & Ruggeri, A. (2018). Active function learning, In Kalish, C., Rau, M., Zhu, J., & Rogers, T. (Eds.), *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*, (pp. 578-583), Madison, WI: Cognitive Science Society.

As the first author, I was responsible for 65% of the work involved in this publication: task elaboration and data collection, statistical analyses, writing the manuscript, responding to reviews, and implementing revisions. Dr. Eric Schulz contributed his expertise in computational modeling by building and testing the models presented in the paper (20% of the work). Prof. Dr. Björn Meder (10%) and Prof. Dr. Azzurra Ruggeri (5%) provided critical feedback to guide the development of the paper.

5.1 Paper 1: How is the hypothesis space represented? Evidence from young children’s active search and predictions in a multiple-cue inference task

To successfully navigate an uncertain world, one has to learn the relationship between cues (e.g., wind speed and atmospheric pressure) and outcomes (e.g., rain). When learning, it is sometimes possible to actively manipulate the cue values, allowing one to test hypotheses about this relationship directly. Across two studies, we investigated how 5- to 7-year-olds selected cue configurations when learning cue-outcome relationships, and what their active search and learning performance revealed about the way they represented these relationships in the hypothesis space. In our task, children had to learn how two cues (color and shape) predicted some monsters’ relative speed, by actively selecting which monster pairs to see running in a race. Based on modeling work with adults, we compared two computational models in their ability to capture children’s

search patterns and learning performance: the *cue-abstraction* model relies on a hierarchical representation that organizes the hypothesis space based on abstracted cue-outcome relationships, and is an efficient way of representing task-relevant information as it supports fast learning and generalization. The *permutation-based* model represents the hypothesis space in terms of the relative speed of individual monsters (i.e., it includes all the possible monster rankings), and is therefore more information-intensive and less efficient. The results of Study 1 (26 5-year-olds, 14 female and 25 6-year-olds, 15 female; predominantly white and fluent in English) provided the first evidence that 5- and 6-year-olds can already use cue-abstraction hypothesis space representations when provided with scaffolding, at a much younger age than previously assumed. However, Study 2 (65 5-year-olds, 33 female, 67 6-year-olds, 33 female and 68 7-year-olds, 33 female; predominantly white and fluent in German) showed that young children were best described by the permutation-based model, and that only 7-year-olds, when provided with memory aids, were best captured by the cue-abstraction model. Overall, our results highlight the guiding role of hypothesis space representations for active search and learning, suggesting that these two phases might trigger different representations, and indicating for the first time a developmental shift in how children represent the hypothesis space. This points to changes in the structure and stability of the hypothesis space as potential sources of developmental change in active learning strategies, and also highlights the role of executive functions such as working memory in constraining the development of hypothesis space representations.

5.2 Paper 2: Active function learning

How do people actively explore to learn about functional relationships, that is, how continuous inputs map onto continuous outputs? We introduce a novel paradigm in the form of a card game to investigate information search in continuous, multi-feature function learning scenarios. In the card game, each card depicted a different monster, along

with three feature values (friendly, cheeky and funny) and a criterion value, which were related according to a linear function. In order to learn the underlying function, participants ($n = 98$ adults, recruited from Amazon MTurk and tested online; $n = 45$ active learners, $n = 53$ passive learners) either actively selected or passively observed information in a set of training cards, before moving on to two prediction tasks which assessed how well they had learned the function. In contrast to other active learning tasks, we found no benefit of active learning over passive learning, with participants in both conditions performing similarly well. Using computational modeling, we developed and compared different variants of more traditional rule-based (linear regression), and non-parametric (Gaussian process regression) learning models, paired with different active sampling strategies, to model participants' active learning behavior. Our results showed that participants' performance was best described by a rule-based model that attempts to efficiently learn linear functions with a focus on high and uncertain outcomes (i.e., upper confidence bound [UCB] sampling). These results suggest that participants adapted well to the linear study environment and adopted a similar approach to exploration-exploitation tasks when choosing what information to look at during each step of their search. This study advances our understanding of how people actively search for information to learn about functional relations in the environment and points to potential transfers of sampling strategies between certain kinds of tasks. Changes in the extent of this transfer, as well as in the active sampling strategies themselves, may potentially contribute to developmental changes in active learning strategies.

Supplement 1 presents a follow-up study which investigates the effects of function type and search horizon (i.e., the amount of information available to learners) on active learning behavior and the relative performance of active versus passive learners. In this follow-up study, participants' active learning behavior was best-described by a non-parametric learning model, which can flexibly learn any type of function rather than only linear functions like the rule-based model we considered, and participants' active

sampling strategies were also consistent with UCB sampling. We also found limited benefits of active over passive learning, which only applied to specific applications of participants' newly acquired knowledge. This second study therefore provides more nuanced insights into active function learning strategies and the relative benefits of active over passive learning.

6 General Discussion

This dissertation investigated how two crucial factors directly related to the implementation of active learning strategies —the hypothesis space and sampling strategies — are implemented in children and adults, and how changes in these factors may lead to documented developmental and task-related differences, as well as how they may interact with certain cognitive skills.

The studies presented in this dissertation are in line with previously identified developmental shifts in active learning strategies between the ages of 5 and 10, and have broadened previous findings to a wider range of tasks. Our results suggest that children’s active learning strategies generally progress from less cognitively complex strategies such as considering one hypothesis at a time to strategies requiring more sophisticated skills such as representing information hierarchically and attending to multiple hypotheses at once. Furthermore, these studies also show that this progression can manifest in counterintuitive ways. For instance, while younger children favor search strategies which focus on one hypothesis at a time in question asking and some causal learning tasks, in tasks like multiple-cue inference, they engage in more information-intensive strategies, which, at first glance, may seem to induce greater cognitive load but which actually incur lower load due to the less complex nature of the hypothesis space, as shown in the first publication (Jones et al., 2021, in press).

Furthermore, this publication confirmed the link between cognitive skills such as executive functions and the ability to implement efficient search strategies, in line with findings from work with adults (Hoffmann, von Helversen, & Rieskamp, 2013; Smith, Patalano, & Jonides, 1998). Although the memory load manipulation in Jones et al. (2021, in press) was imprecise, its effect on children’s hypothesis space and performance strongly suggests working memory, and probably other executive functions, either constrain or are tied to the ability to reason about information in a more abstract, hierarchical manner and therefore implement and learn from more efficient active learning

strategies. While executive functions are likely to have strong links with learning strategies in any kind of task, being crucial cognitive skills, many other cognitive factors are probably also related to information search strategies more generally. For example, as discussed in Supplement 2, metacognition is an important factor in causal learning, and targeting students' metacognitive skills leads to improvements in their self-led science learning (Zepeda, Elizabeth Richey, Ronevich, & Nokes-Malach, 2015). Sensitivity to probabilities is also present from an early age and crucial for strategy adaptiveness, as discussed in the Introduction (Chapter 3) and in Supplement 2, and is likely to be applicable to most kinds of search strategies.

Another key finding was the first evidence that the hypothesis space is not necessarily stable throughout a task, but can change between different phases of a task (here, it changed between the search and prediction phases of the multiple-cue inference task in Jones et al., 2021, in press). This raises a number of important research questions, which are outlined in section 6.1 (Future directions). As a result of exploring how children's hypothesis space is represented in multiple-cue inference tasks, it is also less clear how directly active learning strategies can be mapped onto a learner's hypothesis space. The assumption from work with adults (e.g., Markant & Gureckis, 2012; Markant et al., 2016) and from the question asking literature (e.g., Herwig, 1982; Ruggeri & Feufel, 2015; Ruggeri et al., 2016) was that this mapping should be relatively direct, but our findings from this paper paint a more nuanced picture, as the youngest children (5-year-olds) did not use a hypothesis space structure we considered, although they were still able to perform well, indicating that we were unable to capture their approach to the task.

On one hand, our candidate hypothesis space models were inspired by work with adults (Enkvist et al., 2006; Juslin, Jones, et al., 2003; Juslin, Karlsson, & Olsson, 2008; Juslin, Olsson, & Olsson, 2003) and may therefore not have been suitable to describe young children's strategies. On the other hand, the fact that they did capture older

children's strategies, and that even 5-year-olds were able to reason about hierarchical cue relationships when trained using a forced-choice learning phase, raises the possibility that the ability to represent task-relevant information hierarchically may not have been the only factor to determine whether or not children used such a representation to guide their search. Therefore, it is also possible that the mismatch between our candidate models of the hypothesis space and 5-year-olds' search patterns was caused by 5-year-olds' strategies not mapping onto the kinds of hypothesis space representations we considered, rather than because they were simply unable to create such a representation. The mapping between search strategies and hypothesis space representations may also undergo developmental changes, or drive them.

Another important finding was the fact that sampling strategies may not always depend on the specific task, as participants used UCB (Upper Confidence Bound) sampling in the function learning tasks presented in the second paper (Jones et al., 2018, see also the follow-up study in Supplement 1), echoing sampling strategies from tasks involving exploration-exploitation trade-offs. This indicates that sampling strategies may be conserved across a range of tasks. Although this question has not often been explicitly investigated, this result is consistent with efforts to determine which measures of information utility best describe human information search, which showed that probability gain (wherein information is sought based on how much it increases the probability of finding the correct answer) best described sampling strategies in different kinds of tasks (e.g., experience learning and summary-statistics-based tasks; Nelson, McKenzie, Cottrell, & Sejnowski, 2010). However, note that the tasks used were all related to probabilistic category-learning and therefore did not constitute as broad a range of tasks as would be ideal to truly capture how far sampling strategies might be transferred between tasks. Moreover, these results also stand in contrast to later findings which showed that search-payoff structures determined what kind of information adults preferred (Meder & Nelson, 2012). Therefore, the evidence on the extent to which sampling strategies

are conserved between tasks is mixed, but our findings do suggest that there is some transfer, at least within a certain range, which may have important implications for the development of active learning strategies (see section 6.1).

In addition, Jones et al. (2018) also contributed to a more nuanced perspective on the relative benefits of active versus passive learning. We found no advantage of active over passive learning in this study, which suggested either that there truly was no benefit of this learning condition in our task, or that there was a ceiling effect due to the high number of learning trials which may have masked any benefits of active learning. The first interpretation would be consistent with evidence that the benefits of active learning depend on the specific task (e.g., Enkvist et al., 2006; Henriksson & Enkvist, 2018) and applications of learners' knowledge (Plancher et al., 2013), while the second would suggest that the benefits of active learning may depend on the amount of information the learner has access to, in line with Steyvers et al. (2003).

Supplement 1, which examined the effects of search horizon (i.e., the amount of information available) and function type on active function learning using the same paradigm, showed that active learning and a longer search horizon were only advantageous when participants had to explicitly predict a criterion value, rather than compare two sets of cue values. However, note that improvements in learning outcomes from active learning were not mediated by search horizon, suggesting that the benefits of active learning did not depend on how much information learners could uncover during learning, as was theorized from the results of Jones et al. (2018). Further, this advantage was only present when participants' predictions fell outside the range of values they had trained on, and did not depend on the difficulty of the function to be learned. This suggests that the first interpretation of the results of Jones et al. (2018) was more likely: any improvements in learning outcomes from active learning were only relevant for very specific applications of learners' knowledge. Therefore, any benefits of active learning depend on the specific task and learning goal, and it should not automatically

be considered more effective than passive learning.

Based on this premise, it is interesting to consider the kinds of tasks in which active and passive learning have been compared, as several tasks where active learning was more clearly advantageous could broadly be classified as “explicit” learning tasks, in the sense that they were tasks where the learning goal was to learn facts or items that needed to be explicitly retrieved from memory (e.g., in memorization; Ruggeri, Markant, et al., 2019). In contrast, learning in tasks such as function learning and some kinds of categorization was found to be broadly comparable between learning conditions, or even better with passive learning (Enkvist et al., 2006; Henriksson & Enkvist, 2018; Jones et al., 2018). These kinds of tasks could be considered more “implicit”, in the sense that explicit, verbalizable knowledge of the material to be learned, such as a function, is not required for successful learning. In these cases, it might make sense for active learning to provide limited advantages, as any experience with the material, whether it is under the learner’s control or not, would be enough for learning. However, note that this tentative distinction is not necessarily so clear-cut; a lack of advantage from active learning can also be caused by learners’ use of an inappropriate learning strategy rather than because active learning is inherently not beneficial (e.g., Henriksson & Enkvist, 2018).

6.1 Future directions

Several future research directions arise from the findings of this dissertation, chiefly related to achieving a better understanding of the development of the hypothesis space and sampling strategies, and identifying more precise potential targets for interventions aiming to optimize active learning strategies.

First, it is an open question whether the change in hypothesis space found in Jones et al. (2021, in press) was due to participants’ young age, and would therefore stabilize with time, or if this malleability may continue into adulthood, and may depend on task

difficulty. Further research is needed to clarify these questions. If this is a function of age, changes in the ability to reliably represent task-relevant information in a more abstract fashion may therefore drive the development of information search strategies. If not, it would be important to determine the precise conditions under which such changes in hypothesis space representation occur. For instance, it may be that being able to guide information search directly using the hypothesis space leads to increases in strategy efficiency, perhaps because this allows learners to more effectively reduce uncertainty between specific hypotheses. Future research should seek to clarify these questions, as this would be helpful in identifying precise targets for any interventions seeking to improve active learning strategies.

Furthermore, achieving a more complete understanding of how learners' hypothesis space representations are structured in a wider variety of tasks, both in and out of the classroom, may be helpful in determining whether learners have any misconceptions about the material to be learned which may hinder not only their learning outcomes but also the active learning strategies they use. For instance, the supporting study presented in Supplement 2 has shown that it is crucial to be able to determine whether students hold any misconceptions about the best approaches to different kinds of causal learning tasks, as students may resist applying strategies such as CVS if they do not believe it is the best approach. This would help avoid situations where students misapply strategies due to such misunderstandings, rather than because they have not learned the strategies correctly. Investigating these questions would help develop meta-strategic interventions aimed at resolving these misunderstandings and facilitate efficient and effective learning strategies tailored to each individual learner.

Second, Jones et al. (2018) also showed some level of transfer of sampling strategies between different kinds of tasks. It would be important to determine what tasks in which certain sampling strategies are preferred have in common, as well as how this element of information search changes with age. For example, it is possible that children

may start by over-generalizing certain sampling strategies and gradually become more selective, or the opposite may be true, with children being highly selective at first and then beginning to transfer some sampling strategies to related tasks. Adults' ability to tailor their information sampling to reward structure, when considered with the results of the studies in this dissertation, suggests that they have a variety of such strategies in their toolbox, which they apply according to the task characteristics and their specific goals. This would be in line with the principles of the adaptive toolbox, an idea proposed by Gigerenzer and Todd (1999), but this possibility must be investigated directly.

In addition, another important research direction would be to draw a clearer link between specific cognitive skills and the quality and development of active learning strategies. For instance, it would be helpful to determine the extent to which skills and characteristics such as categorization, metacognition and socioeconomic status (explored in Jones, Swaboda, & Ruggeri, 2020), not just executive functions, relate to the development and quality of active learning strategies. Targeting these skills could be a viable way to improve children's learning strategies, instead of trying to optimize the strategies directly.

On a related note, the circumstances under which active learning is preferable to passive learning, as well as the reasons for it being better or worse than passive learning, also need to be explored in more depth. This dissertation points to the task context and, in some situations, cognitive skills, as two factors which are likely to impact the relative benefits of these two learning conditions. The literature on active learning has also highlighted the importance of applying the right kind of learning strategy to the right kind of task, a skill which may also undergo developmental changes, and which even adults still occasionally struggle with, as shown by Henriksson and Enkvist (2018). This may also be a potential target for interventions, either directly or, perhaps, through instruction in metacognition.

6.2 Conclusion

In sum, the results of the studies presented in this dissertation have provided important insights into the mechanics of how active learning strategies are implemented at the computational level and which factors may undergo or drive developmental changes, and in doing so have raised several additional research questions that should be pursued in order to achieve a more complete understanding of active learning. As such, this dissertation takes some of the first steps towards building up a scientific framework that could potentially guide instructors and educational app creators in designing interventions to properly harness and enhance these cognitive and computational processes to boost children's learning.

7 References

- Betsch, T., Lang, A., Lehmann, A., & Axmann, J. M. (2014). Utilizing probabilities as decision weights in closed and open information boards: A comparison of children and adults. *Acta Psychologica, 153*, 74–86. doi: 10.1016/j.actpsy.2014.09.008
- Boekaerts, M. (1997). Self-regulated learning: A new concept embraced by researchers, policy makers, educators, teachers, and students. *Learning and instruction, 7*(2), 161–186. doi: 10.1016/S0959-4752(96)00015-1
- Bonawitz, E. B., van Schijndel, T. J. P., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology, 64*(4), 215–234. doi: 10.1016/j.cogpsych.2011.12.002
- Bonwell, C. C., & Eison, J. A. (1991). *Active learning: Creating excitement in the classroom. 1991 ashe-eric higher education reports*. ERIC.
- Bramley, N. R., Dayan, P., Griffiths, T. L., & Lagnado, D. A. (2016). Formalizing Neurath’s Ship: Approximate Algorithms for Online Causal Learning. , 1–86. doi: 10.1037/rev0000061
- Bruner, J. S. (1961). *The act of discovery*. Harvard Education Publishing Group.
- Chase, C. C., & Klahr, D. (2017). Invention Versus Direct Instruction: For Some Content, It’s a Tie. *Journal of Science Education and Technology, 26*(6), 582–596. doi: 10.1007/s10956-017-9700-6
- Chen, Z., & Klahr, D. (1999). All other things being equal: Children’s acquisition of the control of variables strategy. *Child Development, 70*(5), 1098–1120.
- Chi, M. T. (2009). Active-Constructive-Interactive: A Conceptual Framework for Differentiating Learning Activities. *Topics in Cognitive Science, 1*(1), 73–105. doi: 10.1111/j.1756-8765.2008.01005.x
- Chi, M. T., & Wylie, R. (2014). The ICAP Framework: Linking Cognitive Engagement to Active Learning Outcomes. *Educational Psychologist, 49*(4), 219–243. doi:

10.1080/00461520.2014.965823

- Coenen, A., Ruggeri, A., Bramley, N. R., & Gureckis, T. M. (2019). Testing one or multiple: How beliefs about sparsity affect causal experimentation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *45*(11), 1923–1941. doi: 10.1037/xlm0000680
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011, sep). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, *120*(3), 341–349. doi: 10.1016/J.COGNITION.2011.03.003
- Davidson, D. (1991a). Children's Decision-Making Examined with an Information-Board Procedure. *Cognitive Development*, *6*(1), 77–90. doi: 10.1016/0885-2014(91)90007-Z
- Davidson, D. (1991b). Developmental differences in children's search of predecisional information. *Journal of Experimental Child Psychology*, *52*, 239–255. doi: 10.1016/0022-0965(91)90061-V
- Davidson, D. (1996). The effects of decision characteristics on children's selective search of predecisional information. *Acta Psychologica*, *92*(3), 263–281. doi: 10.1016/0001-6918(95)00014-3
- De Simone, C., & Ruggeri, A. (2020). What is a good question asker better at? From unsystematic generalization to adult-like selectivity across childhood. *Preprint*, 1–27. doi: 10.31219/osf.io/6jupw
- Dean, D. J., & Kuhn, D. (2007). Direct instruction vs. discovery: the long view. *Science Education*, *91*(1), 36–74. doi: 10.1002/sce.20194
- Denison, S., Reed, C., & Xu, F. (2013). The emergence of probabilistic reasoning in very young infants: Evidence from 4.5- and 6-month-olds. *Developmental psychology*, *49*(2), 243. doi: 10.1037/a0028278
- Denison, S., & Xu, F. (2014). The origins of probabilistic inference in human infants. *Cognition*, *130*(3), 335–347. doi: 10.1016/j.cognition.2013.12.001

- De Simone, C., & Ruggeri, A. (2019). What is a good question asker better at? from no generalization, to overgeneralization, to adults-like selectivity across childhood. In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st Annual Conference of the Cognitive Science Society* (pp. 1613–1619). Montreal, QB: Cognitive Science Society. doi: 10.31219/osf.io/6jupw
- Eaves, J. B. S., & Shafto, P. (2012). Unifying pedagogical reasoning and epistemic trust. *Advances in Child Development and Behavior*, 43, 295–319. doi: 10.1016/B978-0-12-397919-3.00011-3
- Edelsbrunner, P. A., Schalk, L., Schumacher, R., & Stern, E. (2018). Variable control and conceptual change: A large-scale quantitative study in elementary school. *Learning and Individual Differences*, 66(August), 38–53. doi: 10.1016/j.lindif.2018.02.003
- Enkvist, T., Newell, B., Juslin, P., & Olsson, H. (2006). On the role of causal intervention in multiple-cue judgment: Positive and negative effects on learning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 32(1), 163–179. doi: 10.1037/0278-7393.32.1.163
- Feldman, A., & Acredolo, L. (1979). The effect of active versus passive exploration on memory for spatial location in children. *Child Development*, 698–704. doi: 10.2307/1128935
- Fontanari, L., Gonzalez, M., Vallortigara, G., & Girotto, V. (2014). Probabilistic cognition in two indigenous mayan groups. *Proceedings of the National Academy of Sciences*, 111(48), 17075–17080. doi: 10.1073/pnas.1410583111
- Freeman, S., Eddy, S. L., McDonough, M., Smith, M. K., Okoroafor, N., Jordt, H., & Wenderoth, M. P. (2014). Active learning increases student performance in science, engineering, and mathematics. *Proceedings of the National Academy of Sciences*, 111(23), 8410–8415. doi: 10.1073/pnas.1319030111
- Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: The adaptive tool-

- box. In G. Gigerenzer, P. M. Todd, & The ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 3–34). Oxford University Press.
- Giroto, V., & Gonzalez, M. (2008). Children’s understanding of posterior probability. *Cognition*, *106*(1), 325–344. doi: 10.1016/j.cognition.2007.02.005
- Gonzalez, M., & Giroto, V. (2011). Combinatorics and probability: Six-to ten-year-olds reliably predict whether a relation will occur. *Cognition*, *120*(3), 372–379. doi: 10.1016/j.cognition.2010.10.006
- Gopnik, A., Griffiths, T. L., & Lucas, C. G. (2015). When Younger Learners Can Be Better (or at Least More Open-Minded) Than Older Ones. *Current Directions in Psychological Science*, *24*(2), 87–92. doi: 10.1177/0963721414556653
- Gureckis, T. M., & Markant, D. (2009). Active Learning Strategies in a Spatial Concept Learning Game. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 31st Annual Meeting of the Cognitive Science Society* (pp. 3145–3150). Austin, TX: Cognitive Science Society.
- Gureckis, T. M., & Markant, D. B. (2012). Self-Directed Learning: A Cognitive and Computational Perspective. *Perspectives on Psychological Science*, *7*(5), 464–481. doi: 10.1177/1745691612454304
- Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences*, *107*(20), 9066–9071. doi: 10.1073/pnas.1003095107
- Harman, K. L., Humphrey, G. K., & Goodale, M. A. (1999). Active manual control of object views facilitates visual recognition. *Current Biology*, *9*, 1315–1318. doi: 10.1016/S0960-9822(00)80053-6
- Hendrickson, A. T., Navarro, D. J., & Perfors, A. (2016). Sensitivity to hypothesis size during information search. *Decision*, *3*(1), 62–80. doi: 10.1037/dec0000039
- Henriksson, M. P., & Enkvist, T. (2018). Learning from observation, feedback, and intervention in linear and non-linear task environments. *Quarterly Journal of Ex-*

- perimental Psychology*, 71(2), 545–561. doi: 10.1080/17470218.2016.1263998
- Herwig, J. E. (1982). Effects of age, stimuli, and category recognition factors in children’s inquiry behavior. *Journal of Experimental Child Psychology*, 33(2), 196–206. doi: 10.1016/0022-0965(82)90015-7
- Hirsh-Pasek, K., Zosh, J. M., Golinkoff, R. M., Gray, J. H., Robb, M. B., & Kaufman, J. (2015). Putting Education in “Educational” Apps: Lessons From the Science of Learning. *Psychological Science in the Public Interest*, 16(1), 3–34. doi: 10.1177/1529100615569721
- Hoffmann, J. A., von Helversen, B., & Rieskamp, J. (2013). Deliberation’s Blindsight: How Cognitive Load Can Improve Judgements. *Psychological Science*, 24(6), 869–879. doi: 10.1177/0956797612463581
- Jirout, J., & Klahr, D. (2012). Children’s scientific curiosity: In search of an operational definition of an elusive concept. *Developmental Review*, 32(2), 125–160. doi: 10.1016/j.dr.2012.04.002
- Jones, A., Markant, D. B., Pachur, T., Gopnik, A., & Ruggeri, A. (2021, in press). How is the hypothesis space represented? evidence from young children’s active search and predictions in a multiple-cue inference task. *Developmental Psychology*. (Preprint, Paper 1 of this dissertation) doi: 10.17605/OSF.IO/D75BM
- Jones, A., Schulz, E., Meder, B., & Ruggeri, A. (2018). Active function learning. In C. Kalish, M. Rau, J. Zhu, & T. Rogers (Eds.), *Proceedings of the 40th Annual Meeting of the Cognitive Science Society* (pp. 578–583). Madison, WI: Cognitive Science Society. (Paper 2 of this dissertation)
- Jones, A., Swaboda, N., & Ruggeri, A. (2020). Developmental changes in question-asking. In L. P. Butler, S. Ronfard, & C. K. H. (Eds.), *The Questioning Child: Insights from Psychology and Education* (p. 118-143). Cambridge University Press. doi: 10.1017/9781108553803.007
- Juslin, P., Jones, S., Olsson, H., & Winman, A. (2003). Cue abstraction and exem-

- plar memory in categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(5), 924–941. doi: 10.1037/0278-7393.29.5.924
- Juslin, P., Karlsson, L., & Olsson, H. (2008). Information integration in multiple cue judgment: A division of labor hypothesis. *Cognition*, 106(1), 259–298. doi: 10.1016/j.cognition.2007.02.003
- Juslin, P., Olsson, H., & Olsson, A.-C. (2003). Exemplar effects in categorization and multiple-Cue judgment. *Journal of Experimental Psychology: General*, 132(1), 133–156. doi: 10.1037/0096-3445.132.1.133
- Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T.-y., & Camerer, C. F. (2009). The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science*, 20(8), 963–973. doi: 10.1111/j.1467-9280.2009.02402.x
- Kidd, C., & Hayden, B. Y. (2015). The Psychology and Neuroscience of Curiosity. *Neuron*, 88(3), 449–460. doi: 10.1016/j.neuron.2015.09.010
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLoS One*, 7(5), e36399. doi: 10.1371/journal.pone.0036399
- Kinney, D. K., & Kagan, J. (1976). Infant attention to auditory discrepancy. *Child Development*, 155–164. doi: 10.2307/1128294
- Kirschner, P., Sweller, J., & Clark, R. E. (2006). Why unguided learning does not work: An analysis of the failure of discovery learning, problem-based learning, experiential learning and inquiry-based learning. *Educational Psychologist*, 41(2), 75–86. doi: 10.1207/s15326985ep4102_1
- Klahr, D., Fay, A. L., & Dunbar, K. (1993). Heuristics for scientific experimentation: A developmental study. *Cognitive Psychology*, 25, 111–146.
- Klahr, D., & Nigam, M. (2004). The equivalence of learning paths in early science instruction: Effects of direct instruction and discovery learning. *Psychological*

- science*, 15(10), 661–667. doi: 10.1111/j.0956-7976.2004.00737.x
- Kolb, D. A. (2014). *Experiential learning: Experience as the source of learning and development*. FT press.
- Kuhn, D., & Angelev, J. (1976). An experimental study of the development of formal operational thought. *Child Development*, 697–706. doi: 10.2307/1128184
- Kuhn, D., Black, J., Keselman, A., & Kaplan, D. (2000). The development of cognitive skills to support inquiry learning. *Cognition and Instruction*, 18(4), 495–523. doi: 10.1207/S1532690XCI1804_3
- Kuhn, D., & Brannock, J. (1977). Development of the isolation of variables scheme in experimental and “natural experiment” contexts. *Developmental Psychology*, 13(1), 9.
- Kuhn, D., & Dean, D. J. (2005). Is developing scientific thinking all about learning to control variables? *Psychological Science*, 16(11), 866–870. doi: 10.1111/j.1467-9280.2005.01628.x
- Kuhn, D., Garcia-Mila, M., Zohar, A., Andersen, C., White, S. H., Klahr, D., & Carver, S. M. (1995). Strategies of knowledge acquisition. *Monographs of the Society for Research in Child Development*, 60(4), 1–157.
- Kuhn, D., & Phelps, E. (1982). The development of problem-solving strategies. In *Advances in Child Development and Behavior* (Vol. 17, pp. 1–44). Elsevier. doi: 10.1016/S0065-2407(08)60356-0
- Kushnir, T., & Gopnik, A. (2005). Young children infer causal strength from probabilities and interventions. *Psychological Science*, 16(9), 678–683. doi: 10.1111/j.1467-9280.2005.01595.x
- Langsford, S., Hendrickson, A., Perfors, A., & Navarro, D. (2014). People are sensitive to hypothesis sparsity during category discrimination. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 2531–2536). Austin TX: Cognitive Science

Society.

- Legare, C. H., Gelman, S. A., & Wellman, H. M. (2010). Inconsistency with prior knowledge triggers children's causal explanatory reasoning. *Child Development, 81*(3), 929–944. doi: 10.1111/j.1467-8624.2010.01443.x
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin, 116*(1), 75. doi: 10.1037/0033-2909.116.1.75
- Markant, D. B., & Gureckis, T. M. (2012). One piece at a time: Learning complex rules through self-directed sampling. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 725–730). Austin, TX: Cognitive Science Society.
- Markant, D. B., & Gureckis, T. M. (2014). Is it better to select or to receive? Learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General, 143*(1), 94. doi: 10.1037/a0032108
- Markant, D. B., Settles, B., & Gureckis, T. M. (2016). Self-Directed Learning Favors Local, Rather Than Global, Uncertainty. *Cognitive Science, 40*(1), 100–120. doi: 10.1111/cogs.12220
- Mata, R., von Helversen, B., & Rieskamp, J. (2011, mar). When Easy Comes Hard: The Development of Adaptive Strategy Selection. *Child Development, 82*(2), 687–700.
- Matlen, B. J., & Klahr, D. (2013). Sequential effects of high and low instructional guidance on children's acquisition of experimentation skills: Is it all in the timing? *Instructional Science, 41*(3), 621–634. doi: 10.1007/s11251-012-9248-z
- Mayer, R. E. (2004). Should there be a three-strikes rule against pure discovery learning? *American Psychologist, 59*(1), 14. doi: 10.1037/0003-066X.59.1.14
- McComas, J., Dulberg, C., & Latter, J. (1997). Children's memory for locations visited: Importance of movement and choice. *Journal of Motor Behavior, 29*(3), 223–229. doi: 10.1080/00222899709600837

- McKenzie, C. R., Chase, V., Todd, P., & Gigerenzer, G. (2012). Why rare things are precious: The importance of rarity in lay inference. *Ecological rationality: Intelligence in the world*, 309–334.
- Meder, B., & Nelson, J. D. (2012). Information search with situation-specific reward functions. *Judgment and Decision Making*, 7(2), 119–148.
- Mosher, F. A., & Hornsby, J. R. (1966). *On asking questions*. American Psychological Association.
- National Academy of Sciences. (2013). *Next generation science standards: For states, by states*. Washington, DC: National Academies Press.
- Navarro, D. J., & Perfors, A. F. (2011). Hypothesis generation, sparse categories, and the positive test strategy. *Psychological review*, 118(1), 120. doi: 10.1037/a0021110
- Nelson, J. D., McKenzie, C. R. M., Cottrell, G. W., & Sejnowski, T. J. (2010). Experience matters information acquisition optimizes probability gain. *Psychological Science*, 21(7), 960–969. doi: 10.1177/0956797610372637
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101(4), 608.
- Plancher, G., Barra, J., Orriols, E., & Piolino, P. (2013). The influence of action on episodic memory: A virtual reality study. *Quarterly Journal of Experimental Psychology*, 66(5), 895–909. doi: 10.1080/17470218.2012.722657
- Poag, C. K., Cohen, R., & Weatherford, D. L. (1983). Spatial representations of young children: The role of self-versus adult-directed movement and viewing. *Journal of Experimental Child Psychology*, 35(1), 172–179. doi: 10.1016/0022-0965(83)90077-2
- Prince, M. (2004). Does active learning work? A review of the research. *Journal of Engineering Education*, 93(3), 223–231. doi: 10.1002/j.2168-9830.2004.tb00809.x

- Ruggeri, A., & Feufel, M. (2015). How basic-level objects facilitate question-asking in a categorization task. *Frontiers in Psychology, 6*, 918. doi: 10.3389/fpsyg.2015.00918
- Ruggeri, A., & Lombrozo, T. (2015). Children adapt their questions to achieve efficient search. *Cognition, 143*, 203–216. doi: 10.1016/j.cognition.2015.07.004
- Ruggeri, A., Lombrozo, T., Griffiths, T. L., & Xu, F. (2016). Sources of developmental change in the efficiency of information search. *Developmental Psychology, 52*(12), 2159–2173. doi: 10.1037/dev0000240
- Ruggeri, A., Markant, D. B., Gureckis, T. M., Xu, F., Bretzke, M., & Xu, F. (2019, may). Memory enhancements from active control of learning emerge across development. *Cognition, 186*, 1–34. doi: 10.17605/OSF.IO/5FJBT
- Ruggeri, A., Sim, Z. L., & Xu, F. (2017). "Why Is Toma Late to School Again?" Preschoolers Identify the Most Informative Questions. *Developmental Psychology, 53*(9), 1620–1632. doi: 10.1037/dev0000340
- Ruggeri, A., Swaboda, N., Sim, Z. L., & Gopnik, A. (2019). Shake it baby, but only when needed: Preschoolers adapt their exploratory strategies to the information structure of the task. *Cognition, 193*. doi: 10.1016/j.cognition.2019.104013
- Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychological Science, 30*, 1561-1572. doi: 10.1177/0956797619863663
- Schulz, L. (2015). Infants explore the unexpected. *Science, 348*(6230), 42–43. doi: 10.1126/science.aab0582
- Schulz, L. E., & Bonawitz, E. B. (2007). Serious Fun: Preschoolers Engage in More Exploratory Play When Evidence Is Confounded. *Developmental Psychology, 43*(4), 1045–1050. doi: 10.1037/0012-1649.43.4.1045
- Schulz, L. E., Bonawitz, E. B., & Griffiths, T. L. (2007). Can being scared cause tummy aches? naive theories, ambiguous evidence, and preschoolers' causal inferences.

- Developmental Psychology*, 43(5), 1124. doi: 0.1037/0012-1649.43.5.1124
- Smith, E. E., Patalano, A. L., & Jonides, J. (1998, jan). Alternative strategies of categorization. *Cognition*, 65(2-3), 167–196. doi: 10.1016/S0010-0277(97)00043-7
- Sodian, B., Zaitchik, D., & Carey, S. (1991). Young Children’s Differentiation of Hypothetical Beliefs from Evidence. *Child Development*, 62(4), 753–766. doi: 10.1111/j.1467-8624.1991.tb01567.x
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants’ learning and exploration. *Science*, 348, 91 – 94.
- Steffe, L. P., & Gale, J. E. (1995). *Constructivism in education*. Psychology Press.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). *Inferring causal networks from observations and interventions* (Vol. 27) (No. 3). doi: 10.1016/S0364-0213(03)00010-7
- Strand-Cary, M., & Klahr, D. (2008). Developing elementary science skills: Instructional effectiveness and path independence. *Cognitive Development*, 23(4), 488–511. doi: 10.1016/j.cogdev.2008.09.005
- Todd, P. M., & Gigerenzer, G. E. (2012). *Ecological rationality: Intelligence in the world*. Oxford University Press.
- van der Graaf, J., Segers, E., & Verhoeven, L. (2015). Scientific reasoning abilities in kindergarten: dynamic assessment of the control of variables strategy. *Instructional Science*, 43(3), 381–400. doi: 10.1007/s11251-015-9344-y
- van Schijndel, T. J., Jansen, B. R., & Raijmakers, M. E. (2018). Do individual differences in children’s curiosity relate to their inquiry-based learning? *International Journal of Science Education*, 40(9), 996–1015. doi: 10.1080/09500693.2018.1460772
- von Helversen, B., Mata, R., & Olsson, H. (2010). Do children profit from looking beyond looks? From similarity-based to cue abstraction processes in multiple-cue judgment. *Developmental Psychology*, 46(1), 220–229. doi: 10.1037/a0016690

- Voss, J. L., Gonsalves, B. D., Federmeier, K. D., Tranel, D., & Cohen, N. J. (2011). Hippocampal brain-network coordination during volitional exploratory behavior enhances learning. *Nature neuroscience*, *14*(1), 115–120. doi: 10.1038/nn.2693
- Zepeda, C. D., Elizabeth Richey, J., Ronevich, P., & Nokes-Malach, T. J. (2015). Direct instruction of metacognition benefits adolescent science learning, transfer, and motivation: An in vivo study. *Journal of Educational Psychology*, *107*(4), 954–970. doi: 10.1037/edu0000022
- Zimmerman, C. (2007). The development of scientific thinking skills in elementary and middle school. *Developmental Review*, *27*(2), 172–223. doi: 10.1016/J.DR.2006.12.001

8 Appendix

The two published manuscripts at the core of this dissertation are not provided in this dissertation for copyright reasons, but the supporting unpublished manuscripts are provided in the Appendix. The references for each of these are as follows:

Paper 1:

Jones, A., Markant, D. B., Pachur, T., Gopnik, A., & Ruggeri, A. (2021). How is the hypothesis space represented? Evidence from young children's active search and predictions in a multiple-cue inference task. *Developmental Psychology*. Preprint doi: 10.17605/OSF.IO/D75BM

Paper 2:

Jones, A., Schulz, E., Meder, B., & Ruggeri, A. (2018). Active function learning, In Kalish, C., Rau, M., Zhu, J., & Rogers, T. (Eds.), *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*, (pp. 578-583), Madison, WI: Cognitive Science Society.

Supplement 1:

Jones, A., Schulz, E., Meder, B., & Ruggeri, A. (under revision). Learning functions actively. Submitted to *Cognitive Science*.

Supplement 2:

Jones, A., Bramley, N. R., Gureckis, T. M., & Ruggeri, A. (under review). Changing many things at once sometimes makes for a good experiment, and children know that. Submitted to *Developmental Psychology*. Preprint doi: 10.17605/OSF.IO/E3MKC

Learning functions actively

Angela Jones

Max Planck Institute for Human Development

Technical University Munich

Eric Schulz

Max Planck Institute for Biological Cybernetics

Björn Meder

Max Planck Institute for Human Development

Health and Medical University Potsdam

Azzurra Ruggeri

Max Planck Institute for Human Development

Technical University Munich

Author Note

Angela Jones, Max Planck Research Group iSearch, Max Planck Institute for Human Development.

Eric Schulz, Max Planck Research Group Computational Principles of Intelligence, Max Planck Institute for Biological Cybernetics.

Azzurra Ruggeri, Max Planck Research Group iSearch, Max Planck Institute for Human Development and School of Education, Technical University, Munich.

Björn Meder, Max Planck Research Group iSearch, Max Planck Institute for Human Development and Department of Health, Health and Medical University Potsdam.

AJ and ES contributed equally to this work and share first authorship.

Corresponding author: Eric Schulz, E-mail: eric.schulz@tue.mpg.de

Abstract

How do people actively learn functional rules, i.e. a mapping of continuous inputs onto a continuous output? We investigate information search behavior in a multiple-feature function learning task in which participants either actively select or passively receive observations. We find that participants benefit from actively selecting information, in particular in their function extrapolation performance. By introducing and comparing different models of active function learning, we find that participants are best described by a non-parametric function learning model that learns about both the underlying function and inputs that are likely to produce high outputs. These results enrich our understanding of active function learning in complex domains.

Keywords: active learning; function learning; self-directed learning; search

Introduction

In every day life, we often have to learn functional relationships between different variables. How far can I drive with my new electric vehicle when the battery is fully charged? How much breading do I need for the perfect schnitzel? How many rhetorical questions should I pose to make my introduction compelling?

Traditionally, function learning behavior has been studied in passive information-processing paradigms. In these paradigms, participants are sequentially confronted with continuous inputs, for example the length of a horizontal line, followed by a continuous response, such as the length of another horizontal or vertical line (Carroll, 1963; DeLosh, Busemeyer, & McDaniel, 1997; Kalish, 2013). Often these inputs and outputs represent concrete, meaningful variables such as the amount of a chemical substance (inputs) and the resulting amount of arousal in test subjects (outputs; DeLosh et al., 1997; McDaniel, Dimperio, Griego, & Busemeyer, 2009). Participants' task is to learn the underlying function relating inputs to outputs. Learning success can be tested, for instance, by asking participants to make predictions about the outcome variable given previously unobserved input values (i.e., function extrapolation). These experiments have focused on *passive* function learning, where the provided inputs are either randomly determined or selected by the researcher. However, we often actively decide for which inputs we want to observe the outcome in the real world. For instance, to learn about how far one can drive an electric vehicle with a full charge, one could measure the maximum distance covered when driving at different speeds. How can and should an agent actively learn about functional relations among continuous variables? And what models describe human active function learning best?

In this paper, we implement a multiple-feature function learning task to investigate how adult participants actively select inputs for which they want to observe the resulting output. Our behavioral results show that people's understanding of the underlying

function is more accurate when learning actively compared to passively observing randomly selected inputs and corresponding output. The advantage of active over passive learning is particularly pronounced when participants have to make judgments about new inputs (i.e. extrapolation judgements). To better characterize participants' search behavior, we evaluate several combinations of function learning models and active sampling strategies. The best-performing model is a Gaussian Process function learning model combined with an Upper Confidence Bound sampling strategy. This indicates that participants learn functions in a flexible way and can adapt to different underlying functional rules instead of assuming only one particular rule (e.g., a linear function). Moreover, the fact that this model fits best when combined with an Upper Confidence Bound sampling strategy suggests that participants care about both learning the function and finding inputs that produce high outputs.

Function learning

Studies on function learning usually present participants with several input-output pairs (e.g., two bars of different lengths), and then test their learning of the underlying function by asking them to infer the output for inputs that have not been observed before (e.g., to predict the length of a second bar, given the first), either included in the range of the training inputs (*interpolation* task; e.g., the length of the first bar is very similar to one previously observed) or outside the range of training inputs (*extrapolation* task; e.g., the length of the first bar is different from any previously observed).

Studies using interpolation tasks have shown that linear, increasing functions are easier to learn than non-linear, decreasing functions (Brehmer, 1974; Brehmer, Alm, & Warg, 1985; Byun, 1996; McDaniel & Busemeyer, 2005). Studies using extrapolation tasks (DeLosh et al., 1997; McDaniel & Busemeyer, 2005) have demonstrated that participants tend to extrapolate in a linear fashion (Kalish, Lewandowsky, & Kruschke, 2004; Kwantes & Neal, 2006), even when the underlying function is nonlinear (DeLosh et al., 1997).

However, people are capable of non-linear extrapolation (Busemeyer, Byun, Delosh, & McDaniel, 1997), for example when the underlying function is quadratic (Byun, 1996) or cyclical (Bott & Heit, 2004), although the latter case is subject to debate (Kalish, 2013). They therefore have a strong linear bias when learning functional relationships, but remain generally flexible learners, able to adapt to the type of function being learned.

Different theories have been developed to explain these findings and account for human function learning. The most prominent are similarity-based and rule-based theories. *Similarity-based* theories (e.g., Busemeyer et al., 1997; DeLosh et al., 1997) assume that people associate similar inputs with similar outputs, without learning an explicit representation of the underlying function. Similarity-based theories successfully capture some aspects of the observed performance, for instance that some functions are easier to learn than others. However, they fail to explain participants' systematic extrapolation patterns.

Rule-based theories (Carroll, 1963; Koh & Meyer, 1991) assume that participants learn explicit parametric representations, for example linear or power-law functions. Rule-based theories of function learning can successfully predict linear function extrapolation performance, for example by simply assuming that participants learn linear rules. However, they fail to explain that some rules are more difficult to interpolate than others (McDaniel & Busemeyer, 2005).

Hybrid models of function learning contain a similarity-based learning process that acts on explicitly-represented rules. They assume similarity-based interpolation, but extrapolate using simple linear models (Bott & Heit, 2004; Busemeyer et al., 1997; McDaniel & Busemeyer, 2005). Some hybrid models are able to capture both extrapolation and interpolation patterns (McDaniel et al., 2009), such as the EXAM (DeLosh et al., 1997; McDaniel & Busemeyer, 2005) and POLE models (Kalish et al., 2004). For instance, EXAM has been shown to capture participants' linear bias, interpolation and extrapolation performance, but does not account for participants' ability to extrapolate non-linear

functions (Bott & Heit, 2004; McDaniel et al., 2009). POLE does not always capture interpolation or extrapolation performance as well as EXAM, but it does account for the phenomenon of knowledge partitioning (McDaniel et al., 2009).

A related model has been proposed by Griffiths, Lucas, Williams, and Kalish (2009), who have put forward a rational theory of function learning based on Gaussian Process regression. Gaussian Process (GP) regression is a non-parametric method to perform Bayesian regression. Moreover, GP regression exhibits an inherent mathematical duality that makes it compatible with both a rule-based and a similarity-based account of function learning. Gaussian Processes generate predictions based on the similarity between different input values as expressed through a kernel, reminiscent of similarity-based models, and every kernel can be considered the result of performing a Bayesian regression, echoing rule-based models, as each kernel corresponds to a particular prior over functions. Lucas, Griffiths, Williams, and Kalish (2015) and Schulz, Tenenbaum, Duvenaud, Speekenbrink, and Gershman (2017) showed that GP regression can account for a wide range of human interpolation and extrapolation patterns.

Active learning

In the past years, a strong interest in human information search and active learning has emerged, with several studies finding beneficial effects of active compared to passive learning (see Coenen, Nelson, & Gureckis, 2018, for a review). For instance, Lagnado and Sloman (2004) found that learners who were given the opportunity to actively intervene on a causal system made more accurate inferences than passive learners who could not freely decide which information to obtain (see also Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). In category learning, Markant and Gureckis (2014) found that active learners sampled more along the line of the category boundaries, thereby selecting more informative inputs, which in turn increased their categorization performance. Furthermore, recent studies have demonstrated that active control of the study experience leads to enhanced

recognition memory in both children and adults (Markant, Ruggeri, Gureckis, & Xu, 2016; Ruggeri, Markant, Gureckis, Bretzke, & Xu, 2019), compared to conditions lacking this control, and that this benefit persists over time.

However, studies investigating whether active learning is beneficial in multiple-cue learning tasks, which are related to function learning, are less clear (Enkvist, Newell, Juslin, & Olsson, 2006; Osman & Speekenbrink, 2012). Active learning led to learning enhancements in multiple-cue learning but not when the cues were binary (Enkvist et al., 2006) and was no better or worse than passive observation in dynamic environments (Osman & Speekenbrink, 2012), suggesting that the benefits of active learning may depend on the type and context of tasks. Whether the opportunity to learn functions actively results in performance enhancements is an open question.

A critical question discussed in research on active learning is how to define the *usefulness* of pieces of information (see Nelson, 2005; Settles, 2010, for reviews). Different formal measures have been put forward, with the most prominent ones including the reduction in uncertainty measured via Shannon (1948) entropy (Lindley et al., 1956), the increase in the probability of making a correct classification decision (Nelson, McKenzie, Cottrell, & Sejnowski, 2010), and obtaining information for improving payoffs (Meder & Nelson, 2012; Wu, Schulz, Speekenbrink, Nelson, & Meder, 2018). Crupi, Nelson, Meder, Cevolani, and Tentori (2018) demonstrated that several of these measures can be unified into a coherent mathematical framework, thereby connecting formerly competing models of the value of information.

It is still unclear which measure best accounts for how human learners select information. For instance, probability gain consistently best described human search decisions in experienced-based category learning, where the goal is to maximize overall classification accuracy (Meder & Nelson, 2012; Nelson et al., 2010). In other tasks, however, information gain (expected reduction in Shannon entropy) is a better predictor for human search behavior (Bramley, Lagnado, & Speekenbrink, 2015; Markant &

Gureckis, 2012; Meder, Nelson, Jones, & Ruggeri, 2019; Nelson, Divjak, Gudmundsdottir, Martignon, & Meder, 2014). Moreover, search behavior can vary depending on how information about the structure of the environment is communicated (Nelson et al., 2010; Wu, Meder, Filimon, & Nelson, 2017). These findings suggest that there might not be one single measure of usefulness that can account for behavior across all paradigms. Generally, it is still debated which measure of usefulness best describes active learning behavior in more complex domains such as function learning, category learning and causal learning, which require combining a model of learning and a sampling strategy for evaluating and selecting queries (Bramley et al., 2015; Wu et al., 2018).

The present study: Active function learning

In this paper, we investigate the impact of active control over the function learning process on performance. To do that, we propose an experimental and theoretical framework for studying function learning that marries research on human function learning with recent advances in psychological theories of active learning.

Next, we describe the paradigm we developed to investigate active function learning. We then report analyses of the behavioral data, complemented by a computational analysis of participants' learning and search behavior, in which we compare different models of active function learning.

Experiment

Participants

Participants were 720 adults (mean age=36.34, $SD = 10$, 294 females), recruited from Amazon Mechanical Turk. Average task duration was 11.83 minutes ($SD = 10.71$). Participants received a participation fee of \$2.00 and a bonus of up to \$1.40 (mean bonus=\$0.97, SD =\$0.23). Study approval was obtained from the Max Planck Institute Ethical Review Board and participants gave informed consent prior to participating.

Materials and Procedure

Participants played a browser-based card game, in which each card showed a different monster with values for its three features (“friendly,” “cheeky,” and “funny”; see Figure 1). The instructed goal was to learn to predict the number of “magic fruits” monsters picked (criterion), based on their feature values (inputs). Participants were told they would receive a basic participation fee and a performance-dependent bonus.



Figure 1. Screenshot of the multiple-feature function learning task (linear condition). Participants had to learn the relationship between the monsters’ feature values (“friendly,” “cheeky,” and “funny”) and the criterion (number of fruits picked, shown in the top right corner of selected cards). In this example, at this point in the game, the criterion value has been observed for five monster cards, each with a unique feature combination; the criterion values of the remaining cards are unknown.

Methods and design

Participants had to learn the underlying function from sequentially obtaining information on the criterion value for different monsters’ feature values. The learning phase card set consisted of 27 cards generated by factorially combining all feature values between 2 and 4, such that participants observed only a restricted range of the function. All 27

cards were initially displayed with the feature values visible and the criterion value hidden (Figure 1).

Participants were randomly assigned to one of $2 \times 2 \times 5$ between-subject conditions, where we manipulated how people learned about the function (*learning type*, i.e. active or passive), the function underlying the relationship between the monsters' feature values and the criterion (*function type*, i.e. linear or quadratic), and the amount of information participants received during learning (*number of observations*, between 0 and 27).

Learning type. In the active learning condition, participants could choose for which cards to observe the criterion value. Participants in the passive learning condition had to reveal the criterion value of randomly selected cards, one at a time, until the learning horizon was exhausted. Thus, participants in both conditions received the same number of data points, but while active learners could freely decide which data to observe, passive learners received randomly selected data points. Once revealed, the criterion value remained visible throughout the learning phase (Figure 1).

Function type. To test how a possible advantage of active learning might depend on the complexity of the underlying function participants were assigned to either a linear or a quadratic function. The linear function was

$$y = f(x) = 6x_1 + 3x_2 + x_3 - 10, \quad (1)$$

where y is the criterion value and x_1 , x_2 and x_3 are the feature values. The weights for each feature were decaying, to ensure that participants had to attend to all features to achieve good performance and could not easily use simpler strategies, such as tallying.

The quadratic function was

$$y = f(x) = -x_1^2 + 3x_2 + x_3 + 21. \quad (2)$$

We set the weights of the different features such that the range of output values was

similar¹ to that experienced by participants in the linear function condition. For all participants, the features were randomly assigned to x_1 , x_2 or x_3 .

Number of observations. To test how the amount of learning data impacts participants’ function learning under passive vs. active learning regimes, we varied the length of the learning horizon. Participants observed either 0, 1, 5, 22, or 27 input-output pairs (cards) during the learning phase. The group with 0 observations was added to assess how participants would perform when they were not given the chance to gain any information about the underlying function.

Test phase

The test phase consisted of two tasks: a *criterion estimation* task and a *pair comparison* task (order counterbalanced across participants). No feedback was given during the test phases; the final bonus was determined based on participants’ overall performance in the test phase (see below).

In the criterion estimation task, participants had to infer the criterion of a given monster (card) from its feature profile. This task included three types of trials: five recall trials, five interpolation trials, and eight extrapolation trials (18 cards in total; task order was randomized block-wise across participants). In the *recall* trials, the cards presented new monsters but with feature profiles for which participants had already observed the criterion in the learning phase.² In the *interpolation* trials, the cards presented new monsters with feature profiles corresponding to the five cards that had not been observed during the learning phase.³ In the *extrapolation* trials, the cards showed new monsters with feature values of 1 or 5, representing a part of the function space that participants had not

¹The range of outputs for the learning set was 10-30 for the linear function and 13-33 for the quadratic function. The outputs for the extrapolation trials were the same values for both conditions and varied between 0 and 40.

²Note that this means there were no recall trials for the group who only observed one card or no cards at all.

³Note that this means there were no interpolation trials for the group who observed all 27 cards during the learning phase.

been trained on during the learning phase.

For each card, participants were asked to provide their criterion estimates by moving a slider horizontally between 0 and 40 (in increments of 1) until it reached the desired criterion value. Estimates within 5 of the true criterion value were rewarded with \$0.06; estimates within 10 were rewarded with \$0.04; estimates within 20 were rewarded with \$0.02; estimates further than 20 away from the criterion were not rewarded.

In the *pair comparison* task, participants were shown eight card pairs whose feature values ranged between 1 and 5, such that these profiles contained both known and unknown feature values. For each pair, they had to decide which monster had gathered more fruits. This task assessed how well participants could judge the relative weights of each feature in the function they had to learn. For three of these trials, the card pairs were assembled such that one of the three features differed between cards, while the values for the other two features were held constant. For the other five trials, card pairs were assembled so that the value of the first, second or last feature outweighed the combined value of the two other features on each card, so that this feature was the main determinant of the number of fruits collected. Every correct selection was awarded with \$0.04.

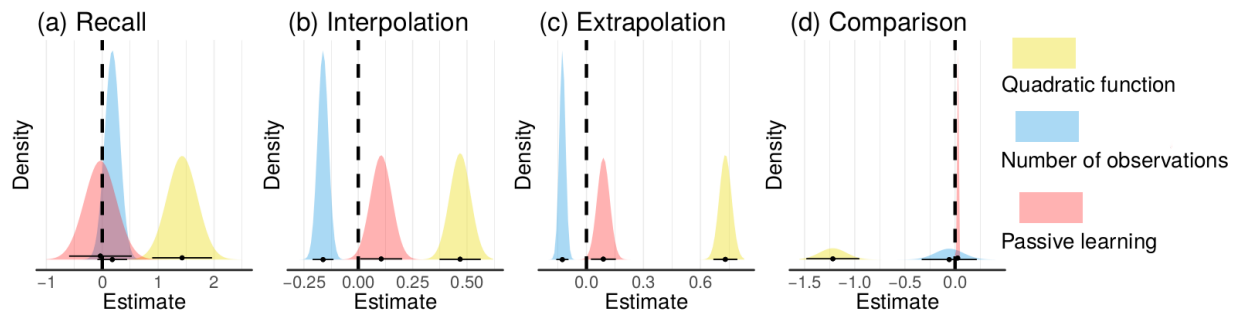


Figure 2. Posterior effects of conditions onto participants' performance. Since performance was measured by participants' absolute error, larger estimates indicate worse performance. (a) Effect on absolute recall error. (b) Effect on absolute interpolation error. (c) Effect on absolute extrapolation error. (d) Effect on pair comparison tasks. Distributions show posterior densities of effects when standardized regression estimates were entered into a Bayesian hierarchical model. Black dots indicate the posterior mean and error bars show the 95% highest posterior density interval.

Behavioral results

We calculated the effect of each manipulation by performing Bayesian multi-level regressions of the conditions' main effects⁴ onto participants' absolute errors in the criterion estimation task (see Appendix for details and additional analyses using maximal random effects structures). For the learning type condition, we created an indicator variable that was set to 1 if learning was passive and 0 if it was active. Function type was coded as 1 if it was quadratic and 0 if it was linear. The number of available observations was entered as continuous variable into the regression. We z -standardized this variable to get a standardized estimate of its effect size. All regressions were performed using a random-intercept over participants and tests are reported based on a comparison with models not containing the tested variable (see Appendix for details).

Figure 2 shows the effects of the different manipulations onto participants' performance (their absolute estimation error) for the different tasks included in the test phase (Figure 2a). For each analysis, we report beta, the estimated highest posterior density (HPD) and Bayes' Factor. If the posterior estimate is negative, the absolute error is lower for these conditions. In the recall trials of the criterion estimation task, we found no evidence for either the horizon ($\beta = 0.18$, $HPD_{95} = [-0.09, 0.44]$, $BF = 0.9$) or the learning type ($\beta = 0.03$, $HPD_{95} = [-0.60, 0.53]$, $BF = 0.5$) being beneficial for participants' performance. However, there was a strong effect of function type, with linear functions being easier to recall than quadratic functions ($\beta = 1.43$, $HPD_{95} = [0.89, 1.96]$, $BF > 100$).

In the interpolation trials, participants performed better when given a longer learning horizon ($\beta = -0.16$, $HPD_{95} = [-0.21, -0.11]$, $BF > 100$) and when learning a linear function ($\beta = -0.47$, $HPD_{95} = [-0.56, -0.37]$, $BF > 100$). We also found moderate evidence for an advantage of the active learning condition ($\beta = 0.11$, $HPD_{95} = [0.01, 0.20]$, $BF = 3$).

⁴There was no evidence for interaction effects with all $BF < 1$.

In the extrapolation trials, we found that participants were better given a longer learning horizon ($\beta = -0.13$, $HPD_{95} = [-0.16, -0.09]$, $BF > 100$) and a linear function ($\beta = 0.73$, $HPD_{95} = [0.67, 0.79]$, $BF > 100$). Additionally, we found evidence for an advantage of the active learning condition ($\beta = 0.09$, $HPD_{95} = [0.02, 0.15]$, $BF > 100$).

Since there was also a group of participants who did not observe any outputs before doing the criterion estimation task (labeled as missing value for the condition variable, since 0 observations are neither active nor passive), we also compared those participants to searchers who had actively sampled only 1 card. This showed that participants who had observed only 1 card already performed better than participants who observed no card at all in the interpolation trials ($\beta = -1.24$, $HPD_{95} = [-2.29, -0.20]$, $BF = 4$) but not in the extrapolation trials ($\beta = 0.33$, $HPD_{95} = [-1.06, 1.77]$, $BF = 0.08$). Thus, we found some evidence that even small amounts of information can improve participants' performance.

To assess participants' performance in the pair comparison task, we calculated the number of correct choices per participant and regressed the different conditions onto this number in a Bayesian linear regression without any random effects⁵. The results revealed that participants did not benefit from actively learning the function ($\beta = -0.06$, $HPD_{95} = [-0.33, 0.22]$, $BF = 0.4$), but performed better in the linear than in the quadratic condition ($\beta = -1.22$, $HPD_{95} = [-1.48, -0.95]$, $BF > 100$) and after having observed more cards ($\beta = 0.03$, $HPD_{95} = [0.01, 0.04]$, $BF = 79$).

Modeling active function learning

Active function learning requires an agent to build up a model of the underlying function and to sample the most useful inputs according to their beliefs. Thus, the building blocks for a computational analysis of active function learning are a model of participants' function learning and of their sampling strategies (e.g., to measure the usefulness of their

⁵Note that we obtain the same results if we treat each response individually as in the analyses for participants' criterion estimation performance. Additionally, the results did not change when performing the analysis for the different types of pairs, i.e. pairs where one feature differed and pairs where two features outweighed another feature.

selection, akin to information gain or probability gain), used to match the model’s expectations onto informative actions. We compared two models of function learning, each combined with three different sampling strategies, to see which combination best accounted for participants’ behavior.

Models for function learning

Linear Regression. A linear regression assumes that the outputs at time point t are a linear function of the inputs plus some added noise:

$$y_t = f(x_t) + \epsilon_t = \beta_0 + \sum_{i=1}^k \beta_i x_{t,i} + \epsilon_t, \quad (3)$$

where the noise term ϵ_t follows a normal distribution $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$ with mean 0 and variance σ_ϵ^2 , β_0 is the intercept term and β_i are the slopes for the different features. Within a Bayesian framework, we can compute the posterior distribution over the weights and use this distribution to generate predictions about new observations, given their feature values (see Appendix for details).

Gaussian Process Regression. A GP regression is a non-parametric Bayesian way to model regression problems that can theoretically learn any stationary function by the means of Bayesian inference (Schulz, Speekenbrink, & Krause, 2018). If f is a function over input space \mathcal{X} that maps to real-valued scalar outputs, then this function can be modeled as a random draw from a GP:

$$f \sim \mathcal{GP}(m, k). \quad (4)$$

Here, m is a mean function that is commonly set to 0 to simplify computations. The kernel function k specifies the covariance between outputs.

$$m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})] \tag{5}$$

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E} [(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))]. \tag{6}$$

The kernel function k encodes prior assumptions about the underlying function. A common choice is the *radial basis function* (RBF) kernel to model the underlying functional dependencies:

$$k_{\text{RBF}}(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{\lambda}\right). \tag{7}$$

The length-scale λ governs the amount of correlation between inputs \mathbf{x} and \mathbf{x}' . Importantly, whereas a linear regression makes explicit assumptions about the underlying functional form (i.e., linear), GP regression makes predictions for new observations based on their similarity to previously observed features and their outputs via the the kernel, which only assumes that the underlying function is locally smooth.

Active Sampling Strategies

Both function learning models generate predictions about the expected mean and associated uncertainties of outputs produced by different inputs. However, active function learning also requires a sampling strategy that maps models' predictions onto utilities to guide data selection. We compared three such sampling strategies.

Uncertainty sampling selects at each step the combination of feature values for which the predicted output is most uncertain, i.e., shows the highest predictive posterior standard deviation.

$$a_t(\mathbf{x}) = \arg \max \sigma_{t-1}(\mathbf{x}) \tag{8}$$

This strategy reduces the uncertainty over the input space quickly, and is mathematically related to focusing on the information gain of each observation (Krause, Singh, & Guestrin, 2008).

Mean sampling selects at each step the input values that currently promise to produce the highest output:

$$a_t(\mathbf{x}) = \arg \max \mu_{t-1}(\mathbf{x}) \quad (9)$$

This strategy does not attempt to learn efficiently but rather learns about the function serendipitously by trying to produce high outputs (i.e., here, higher numbers of magic fruit).

Finally, *upper confidence bound sampling* (UCB) tries to both reduce uncertainty and achieve high outcomes by sampling the input that currently shows the highest upper confidence bound

$$a_t(\mathbf{x}) = \arg \max \mu_{t-1}(\mathbf{x}) + \beta \sigma_{t-1}(\mathbf{x}), \quad (10)$$

where β is a free parameter governing the extent to which participants sample uncertain options. UCB sampling will, on average, converge to both high knowledge about the underlying function and sampling the highest possible outcomes. It has been found to describe human behavior well in exploration-exploitation paradigms where a global value function governs outcomes (Schulz, Wu, Ruggeri, & Meder, 2019; Wu et al., 2018).

Model Comparison Results

We combined all of the above-described function learning models and sampling strategies and compared how well they described individual active learners' card choices in the learning phase (see Fig. 3). Since assessing model accuracy requires more than one data

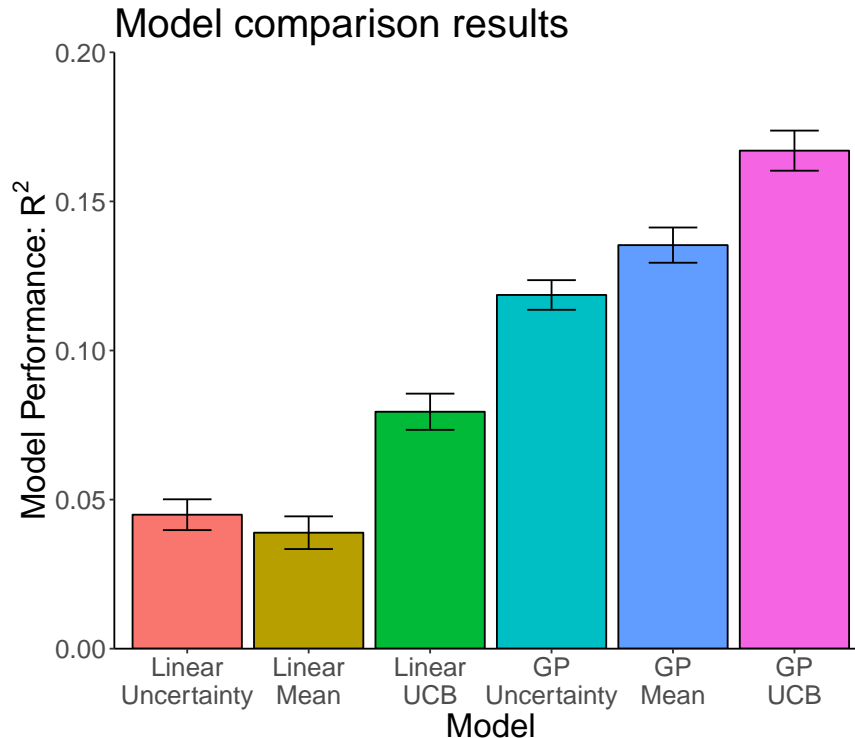


Figure 3. Model comparison results. Average descriptive performance (mean performance over subjects, R^2) for every model and sampling strategy combination. Error bars indicate the standard error of the mean.

point, we only consider learners who observed more than one card (i.e., learning horizon > 1 , $N = 194$).

We used all models to generate means and uncertainties for trial $t + 1$ by feeding them with participants' observations up until trial t , repeating this procedure for every participant over all trials. These means and uncertainties were then converted into utilities by different sampling strategies. Afterwards, we submitted the resulting utilities into a softmax function to convert them into choice probabilities

$$P(x) = \frac{\exp(a_t(\mathbf{x})/\tau)}{\sum_{j=1}^N \exp(a_t(\mathbf{x})/\tau)} \quad (11)$$

where τ is a free temperature parameter estimated for each subject from the data. For each participant, we calculated a model's $\text{AIC}(\mathcal{M}) = -2\log(L(\mathcal{M})) + 2k$, where L is the models

log-likelihood and k the number of free parameters. Afterwards, we standardized model performance using a pseudo- R^2 measure that compared each model to a random baseline (i.e., a model that randomly chooses input combinations (cards) to learn about the function):

$$R^2 = 1 - \frac{\text{AIC}_i}{\text{AIC}_{\text{random}}} \quad (12)$$

Figure 3 shows that the GP function learning model outperformed the linear model for each sampling strategy. The best overall model was a GP regression model combined with a UCB sampling strategy (GP-UCB), which showed an average performance of $R^2 = .17$ and best described 89 participants (i.e., had highest R^2). The second best model was a Gaussian Process regression model combined with a mean sampling strategy, which showed an average performance of $R^2 = .13$ and best described 23 participants. This model performed significantly worse than the GP-UCB model, $t(193) = 9.38$, $p < .001$, $d = 0.67$, $BF > 100$. A Gaussian Process regression model combined with an uncertainty sampling strategy led to an average performance of $R^2 = .12$, and described 60 participants best. This model also performed worse than the GP-UCB model, $t(193) = 6.41$, $p < .001$, $d = 0.46$, $BF > 100$.

The linear regression model combined with an upper confidence bound sampling strategy (Lin-UCB) achieved an average performance of $R^2 = 0.08$, describing 15 participants best overall. This model also performed worse than the GP-UCB model, $t(193) = 11.04$, $p < .001$, $d = 0.79$, $BF > 100$. However, the Lin-UCB model performed better than a linear regression model combined with uncertainty sampling, ($t(193) = 8.16$, $p < .001$, $d = 0.59$, $BF > 100$, which had a mean performance of $R^2 = 0.04$ and described 3 participants best overall. The Lin-UCB model also performed better than the a linear regression model combined with a mean greedy sampling strategy, ($t(193) = 10.41$,

$p < .001$, $d = 0.75$, $BF > 100$, which had a mean performance of $R^2 = .04$ and described 4 participants best overall. Interestingly, the GP-UCB model performed better than the Lin-UCB model even when the underlying function was linear, ($t(85) = 7.97$, $p < .001$, $d = 0.86$, $BF > 100$, indicating that participants applied a Bayesian similarity-based learning strategy even if the underlying function could have been learned by linear rules.

We also analyzed the parameter estimates of the winning GP-UCB model. The mean of the softmax temperature parameter was estimated to be $\hat{\tau} = 0.25$, suggesting that participants' sampling behavior corresponded closely to selecting the highest value option, once they had taken into account both the mean and uncertainty associated with the inputs, and that they did not simply sampled options randomly. The mean estimate for the exploration parameter was $\hat{\beta} = 5.73$, showing that participants valued the reduction of uncertainty positively, trying to learn more about uncertain parts of the underlying function.

Discussion

We investigated participants' function learning behavior and performance in a task where they had to learn about a function relating three continuous features to a continuous criterion. Our behavioral results showed that participants struggled more when having to learn a nonlinear as compared to a linear function. This replicates previous results on human function learning using single features (Brehmer, 1974; Carroll, 1963). Participants' judgments were also more accurate when they could make more observations during the learning phase, in particular when making interpolation and extrapolation judgments. Most importantly, participants benefited from actively learning about the underlying function. This effect was particularly pronounced for function extrapolation in the criterion estimation task. Since extrapolation is known to be a particularly challenging aspect of function learning and has been termed the "sine qua non" of function learning (DeLosh et al., 1997), our findings highlight the advantages of actively learning about functional

relations. Participants in the active learning condition did not, however, show increased performance in the recall and the paired comparison tasks. Because recent results have shown that active learning leads to improved memory of encountered exemplars (Ruggeri et al., 2019), our results suggest that the specific learning goal (e.g., learning a function versus memorizing objects) might mediate the benefits of active learning. Follow-up experiments could further investigate the conditions under which active control over the learning experience benefits participants' recognition memory and functional recall.

By comparing several models of active function learning, we found that a GP regression combined with an upper confidence bound sampling strategy explained participants' behavior best (Lucas et al., 2015; Schulz et al., 2017). This means that participants learned about the underlying function in a flexible and adaptive way; it also shows that participants cared about both reducing uncertainty and finding out which inputs produce high outputs. This finding mirrors previous results obtained in contextual and spatially-correlated bandit tasks (Schulz, Konstantinidis, & Speekenbrink, 2016; Wu et al., 2018). In particular, recent studies showed that participants solve the exploration-exploitation dilemma in reinforcement learning problems in a similar fashion (Schulz, Wu, et al., 2019; Wu et al., 2018). This hints at the possibility of a universal sampling strategy underlying both information search and the search for rewards. Participants may not easily be able to turn off the exploitation part of their sampling strategy as they normally encounter a mix of exploration and exploitation problems in real life (Schulz, Bhui, et al., 2019).

Our results enrich our understanding of active learning in complex domains and pave the way for future studies on active, self-directed function learning. Active function learning is particularly crucial to effectively navigate the world by making accurate inferences and predictions, as many real-world phenomena depend on functional relationships. Mastering this ability can boost learning more generally, especially from a developmental perspective. We know from previous studies that children and adults can

differ in their ability to generalize from past observations and their tendency to seek out uncertainty (Schulz, Wu, et al., 2019). An important question for future research is therefore to identify and precisely characterize the emergence and developmental trajectories of active function learning.

References

- Akaike, H. (1998). Information theory and an extension of the maximum likelihood principle. In *Selected papers of Hirotogu Akaike* (pp. 199–213). Springer.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*, 255–278.
- Bott, L., & Heit, E. (2004). Nonmonotonic extrapolation in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 38–50.
- Bramley, N. R., Lagnado, D. A., & Speekenbrink, M. (2015). Conservative forgetful scholars: How people learn causal structure through sequences of interventions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*, 708–731.
- Brehmer, B. (1974). The effect of cue intercorrelation on interpersonal learning of probabilistic inference tasks. *Organizational Behavior and Human Performance*, *12*, 397–412.
- Brehmer, B., Alm, H., & Warg, L.-E. (1985). Learning and hypothesis testing in probabilistic inference tasks. *Scandinavian Journal of Psychology*, *26*, 305–313.
- Busemeyer, J. R., Byun, E., Delosh, E. L., & McDaniel, M. A. (1997). Learning functional relations based on experience with input-output pairs by humans and artificial neural networks. In K. Lamberts & D. R. Shanks (Eds.), *Studies in cognition. Knowledge, concepts and categories* (pp. 408–437). Cambridge, MA: The MIT Press.
- Byun, E. (1996). *Interaction between prior knowledge and type of nonlinear relationship on function learning*. (Unpublished doctoral dissertation). ProQuest Information & Learning.
- Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, *80*, 1–28.
- Carroll, J. D. (1963). *Functional learning: The learning of continuous functional mappings relating stimulus and response continua*. Educational Testing Service.

- Coenen, A., Nelson, J. D., & Gureckis, T. M. (2018). Asking the right questions about the psychology of human inquiry: Nine open challenges. *Psychonomic Bulletin & Review*, *26*, 1548–1587.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. Hillsdale, NJ: Erlbaum.
- Crupi, V., Nelson, J. D., Meder, B., Cevolani, G., & Tentori, K. (2018). Generalized information theory meets human cognition: Introducing a unified framework to model uncertainty and information search. *Cognitive Science*, *42*, 1410–1456.
- DeLosh, E. L., Bussemeyer, J. R., & McDaniel, M. A. (1997). Extrapolation: The sine qua non for abstraction in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 968–986.
- Enkvist, T., Newell, B., Juslin, P., & Olsson, H. (2006). On the role of causal intervention in multiple-cue judgment: Positive and negative effects on learning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *32*(1), 163–179.
- Griffiths, T. L., Lucas, C., Williams, J., & Kalish, M. L. (2009). Modeling human function learning with gaussian processes. In *Advances in Neural Information Processing Systems* (pp. 553–560).
- Gronau, Q. F., Sarafoglou, A., Matzke, D., Ly, A., Boehm, U., Marsman, M., . . . Steingroever, H. (2017). A tutorial on bridge sampling. *Journal of Mathematical Psychology*, *81*, 80–97.
- Kalish, M. L. (2013). Learning and extrapolating a periodic function. *Memory and Cognition*, *41*(6), 886–896.
- Kalish, M. L., Lewandowsky, S., & Kruschke, J. K. (2004). Population of linear experts: knowledge partitioning and function learning. *Psychological Review*, *111*, 1072–1099.
- Koh, K., & Meyer, D. E. (1991). Function learning: Induction of continuous stimulus-response relations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 811–836.

- Krause, A., Singh, A., & Guestrin, C. (2008). Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, *9*, 235–284.
- Kwantes, P. J., & Neal, A. (2006). Why people underestimate y when extrapolating in linear functions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 1019–1030.
- Lagnado, D. A., & Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 856–876.
- Lindley, D. V., et al. (1956). On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, *27*, 986–1005.
- Lucas, C. G., Griffiths, T. L., Williams, J. J., & Kalish, M. L. (2015). A rational model of function learning. *Psychonomic Bulletin & Review*, *22*, 1193–1215.
- Markant, D., & Gureckis, T. M. (2012). Does the utility of information influence sampling behavior? In *Cogsci*.
- Markant, D., & Gureckis, T. M. (2014). Is it better to select or to receive? learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, *143*, 94–122.
- Markant, D., Ruggeri, A., Gureckis, T. M., & Xu, F. (2016). Enhanced memory as a common effect of active learning. *Mind, Brain, and Education*, *10*(3), 142–152.
- McDaniel, M. A., & Busemeyer, J. R. (2005). The conceptual basis of function learning and extrapolation: Comparison of rule-based and associative-based models. *Psychonomic Bulletin & Review*, *12*, 24–42.
- McDaniel, M. A., Dimperio, E., Griego, J. A., & Busemeyer, J. R. (2009). Predicting transfer performance: A comparison of competing function learning models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 173–195.
- Meder, B., & Nelson, J. D. (2012). Information search with situation-specific reward functions. *Judgment and Decision Making*, *7*, 119–148.

- Meder, B., Nelson, J. D., Jones, M., & Ruggeri, A. (2019). Stepwise versus globally optimal search in children and adults. *Cognition*, *191*, 103965.
- Mullen, K. M., Ardia, D., Gil, D. L., Windover, D., & Cline, J. (2009). DEoptim: An R package for global optimization by differential evolution.
- Nelson, J. D. (2005). Finding useful questions: on Bayesian diagnosticity, probability, impact, and information gain. *Psychological Review*, *112*, 979–999.
- Nelson, J. D., Divjak, B., Gudmundsdottir, G., Martignon, L. F., & Meder, B. (2014). Children’s sequential information search is sensitive to environmental probabilities. *Cognition*, *130*(1), 74–80.
- Nelson, J. D., McKenzie, C. R., Cottrell, G. W., & Sejnowski, T. J. (2010). Experience matters information acquisition optimizes probability gain. *Psychological Science*, *21*(7), 960–969.
- Osman, M., & Speekenbrink, M. (2012). Prediction and control in a dynamic environment. *Frontiers in Psychology*, *3*, 68.
- Rasmussen, C. E., & Nickisch, H. (2010). Gaussian processes for machine learning (gpml) toolbox. *Journal of Machine Learning Research*, *11*(Nov), 3011–3015.
- Rouder, J. N., & Morey, R. D. (2012). Default Bayes Factors for Model Selection in Regression. , *47*, 877–903.
- Ruggeri, A., Markant, D. B., Gureckis, T. M., Bretzke, M., & Xu, F. (2019). Memory enhancements from active control of learning emerge across development. *Cognition*, *186*, 82–94.
- Schulz, E., Bhui, R., Love, B. C., Brier, B., Todd, M. T., & Gershman, S. J. (2019). Structured, uncertainty-driven exploration in real-world consumer choice. *Proceedings of the National Academy of Sciences*, *116*(28), 13903–13908.
- Schulz, E., Konstantinidis, E., & Speekenbrink, M. (2016). Putting bandits into context: How function learning supports decision making. *bioRxiv*, 081091.
- Schulz, E., Speekenbrink, M., & Krause, A. (2018). A tutorial on Gaussian process

- regression: Modelling, exploring, and exploiting functions. *Journal of Mathematical Psychology*, *85*, 1–16.
- Schulz, E., Tenenbaum, J. B., Duvenaud, D., Speekenbrink, M., & Gershman, S. J. (2017). Compositional inductive biases in function learning. *Cognitive Psychology*, *99*, 44 - 79.
- Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychological Science*, *30*, 1561-1572.
- Settles, B. (2010). Active learning literature survey. *University of Wisconsin, Madison*, *52*(55-66), 11.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, *27*, 379–423.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453–489.
- Wu, C. M., Meder, B., Filimon, F., & Nelson, J. D. (2017). Asking better questions: How presentation formats influence information search. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *43*(8), 1274-1297.
- Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, *2*, 915-924.

Appendix A

Statistical tests

We report all statistical tests using both frequentist and Bayesian formats. We present frequentist tests alongside their effect sizes, i.e. Cohen’s d (Cohen’s d ; Cohen, 1988). Bayesian statistics are expressed by their Bayes factors (BFs). A Bayes factor quantifies the likelihood of the data under the alternative hypothesis H_A compared to the likelihood of the data under the null hypothesis H_0 . For example, a BF of 5 indicates that the data are 5 times more likely under H_A than under H_0 ; a BF of 0.2 indicates that the data are 5 times more likely under H_0 than under H_A . We apply the “default” Bayesian t -test as proposed by Rouder and Morey (2012) when comparing two independent groups, using a Jeffreys-Zellner-Siow prior with its scale set to $\sqrt{2}/2$. We approximate the Bayes factor between two different mixed-effects regressions by applying bridge sampling (Gronau et al., 2017). For the Bayesian regression models, we postulate that the β -coefficients for each participant β_i are drawn from a normal distribution

$$\beta_i \sim \mathcal{N}(\mu_\beta, \sigma_\beta^2), \quad (13)$$

estimating the group-level mean μ_β and variance over participants σ_β^2 . We use the following priors on the group-level parameters:

$$\mu_\beta \sim \mathcal{N}(0, 100) \quad (14)$$

$$\sigma_\beta \sim \text{Half-Cauchy}(0, 100). \quad (15)$$

Posterior inference is accomplished by using Hamiltonian Monte Carlo and `brms` package (Bürkner, 2017).

Appendix B

Detailed model implementation

We provide further mathematical details for the two models of active function learning.

Bayesian linear regression

The first function learning model is linear regression. We adopt a Bayesian perspective on linear regression, performing posterior inference over the weights. We assume a Gaussian prior over the weights $p(\mathbf{w}) = \mathcal{N}(0, \Sigma)$ and a Gaussian likelihood $p(y_t | \mathbf{X}_t, \mathbf{w}) = \mathcal{N}(\mathbf{X}_t^\top \mathbf{w}, \sigma_\epsilon^2 \mathbf{I})$. The resulting posterior is

$$\begin{aligned} p(\mathbf{w} | y_t, \mathbf{X}_t) &\propto p(y_t | \mathbf{X}_t, \mathbf{w}) p(\mathbf{w}) \\ &= \mathcal{N} \left(\frac{1}{\sigma_\epsilon^2} \mathbf{A}_t^{-1} \mathbf{X}_t y_t, \mathbf{A}_t^{-1} \right) \end{aligned} \quad (16)$$

where $\mathbf{A}_t = \Sigma^{-1} + \sigma_\epsilon^{-2} \mathbf{X}_t \mathbf{X}_t^\top$. This posterior can be used to generate predictions about different option's means and uncertainties, which can then be fed into different sampling strategies.

Gaussian Process regression

Gaussian process regression assumes that the output y of a function f at input \mathbf{x} can be written as

$$y = f(\mathbf{x}) + \epsilon \quad (17)$$

with $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. In Gaussian process regression, the function $f(\mathbf{x})$ is distributed as a Gaussian process:

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')). \quad (18)$$

A Gaussian process \mathcal{GP} is a distribution over functions and is defined by a *mean* and a *kernel* function. The mean function $m(\mathbf{x})$ reflects the expected function value at input \mathbf{x} : The kernel function k specifies the covariance between outputs.

$$m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})] \quad (19)$$

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))]. \quad (20)$$

Conditional on observed data $\mathcal{D} = \{\mathbf{x}_n, y_n\}_{n=1}^N$, where $y_n \sim \mathcal{N}(f(\mathbf{x}_n), \sigma^2)$ is a noise-corrupted draw from the latent function, the posterior predictive distribution for a new input \mathbf{x}_* is Gaussian with mean and variance given by:

$$\mathbb{E}[f(\mathbf{x}_*)|\mathcal{D}] = \mathbf{k}_*^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y} \quad (21)$$

$$\mathbb{V}[f(\mathbf{x}_*)|\mathcal{D}] = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_*, \quad (22)$$

where $\mathbf{y} = [y_1, \dots, y_N]^\top$, \mathbf{K} is the $N \times N$ matrix of covariances evaluated at each pair of observed inputs, and $\mathbf{k}_* = [k(\mathbf{x}_1, \mathbf{x}_*), \dots, k(\mathbf{x}_N, \mathbf{x}_*)]$ is the covariance between each observed input and the new input \mathbf{x}_* . This posterior distribution can also be used to derive predictions about each option's mean and uncertainties, which can be fed into different sampling strategies.

These predictions depend crucially on the chosen kernel function. The kernel function $k(\mathbf{x}, \mathbf{x}')$ models the dependence between the function values at different input points \mathbf{x} and \mathbf{x}' :

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))] \quad (23)$$

We use a radial basis function kernel, which is defined as

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\lambda^2}\right). \quad (24)$$

The radial basis function provides an expressive kernel to model smooth and stationary functions. The two hyper-parameters λ (called the length-scale) and σ_f^2 (the signal variance) can be varied to increase or reduce the a priori correlation between points and consequentially the variability of the resulting function. We chose those parameters by maximizing the log marginal likelihood. For a GP with hyper-parameters θ , this likelihood is given by:

$$\log p(y|X, \theta) := -\frac{1}{2}y^\top (K + \sigma_n^2 I)^{-1}y - \frac{1}{2} \log |K + \sigma_n^2 I| - \frac{n}{2} \log 2\pi. \quad (25)$$

where the dependence of K on θ is left implicit. We optimize the hyper-parameters using gradient-based optimization as implemented in the `GPML` toolbox (Rasmussen & Nickisch, 2010).

Model comparison

We use both models of learning, the Bayesian linear regression and Gaussian Process regression, to model participants' active learning. We fit the models to a the data a participant has seen time point t and then make predictions about each options mean and choices at $t + 1$. We then feed these means and uncertainties into the different sampling strategies. The resulting utilities are then parsed into a softmax choice rule. We optimize the β of the UCB sampling strategy as well as the temperature parameter τ for each participant using the log-likelihood L . Participant-wise optimization is performed by using differential evolution as implemented in `DEOptim` (Mullen, Ardia, Gil, Windover, & Cline, 2009). The resulting log-likelihood can be used to calculate Akaike's Information Criterion (AIC, Akaike, 1998)

$$\text{AIC}(\mathcal{M}) = -2 \log(L(\mathcal{M})) + 2k, \quad (26)$$

where k indicates the number of optimized parameters (two for any model using UCB sampling, and one otherwise). We standardize the resulting AIC using a pseudo- R^2 measure which compares each model's AIC to a random baseline (without parameters):

$$R^2 = 1 - \frac{\text{AIC}_i}{\text{AIC}_{\text{random}}}. \quad (27)$$

Appendix C

Behavioral results without extreme numbers of observations

Because our data set also contained participants with either 0 or 30 number of observations, we also analyzed our main behavioral effects after excluding these participants. We therefore combined participants' recall, interpolation, and extrapolation performance into one data set of participants' criterion estimation performance. All of the variables were coded as in our main analyses. Additionally, we removed participants with either 0 or 30 observations. We then again estimated a linear Bayesian multi-level model, regressing the number of observations, the function type, and the learning type onto participants' absolute error during the test trials. The results of this analysis showed strong effects for all three of our main manipulations. In particular, participants performed better given a longer learning horizon ($\beta = -0.08$, $HPD_{95} = [-0.10, -0.06]$, $BF > 100$) and a linear function ($\beta = 4.05$, $HPD_{95} = [3.58, 4.51]$, $BF > 100$). Importantly, there was also a strong effect of learning condition, with participants in the active learning condition performing better than participants in the passive condition ($\beta = 0.65$, $HPD_{95} = [0.12, 1.12]$, $BF > 100$).

Next, we also analyzed the behavioral results using maximal linear mixed-effects models (Barr, Levy, Scheepers, & Tily, 2013). Although the overall model comparison suggested that only including a random intercept over participants was enough, it is nonetheless sometimes recommended to keep the comparison maximally. Thus, we repeated the analysis from above, this time entering all of the individual factors as random effects into the null model and comparing them to an alternative model that additionally included the tested factor as a fixed effect as well. This analysis also revealed that all three manipulations had a significant effect onto participants' criterion estimation performance. Specifically, participants performed better given more observations ($\beta = -0.08$, $HPD_{95} = [-0.10, -0.05]$, $BF > 100$), a underlying linear function ($\beta = 4.09$, $HPD_{95} = [3.59, 4.55]$, $BF > 100$) as well as in the active learning condition ($\beta = 0.61$,

$HPD_{95} = [0.16, 1.08], BF > 100$).

**Changing many things at once sometimes makes for a good
experiment, and children know that**

Angela Jones

Max Planck Institute for Human Development, Berlin, Germany & School of
Education, Technical University of Munich, Germany

Neil R. Bramley

Department of Psychology, University of Edinburgh, Scotland

Todd M. Gureckis

Department of Psychology, New York University, USA

Azzurra Ruggeri

Max Planck Institute for Human Development, Berlin & School of Education, Technical
University of Munich, Germany

Correspondence concerning this article should be addressed to Azzurra Ruggeri, MPRG iSearch | Information Search, Ecological and Active Learning research with Children, Max Planck Institute for Human Development, Berlin, Germany, Lentzeallee 94, Berlin, Germany, Phone: + 49 30 82 406 268. E-mail: ruggeri@mpib-berlin.mpg.de

The authors have no conflict of interest to declare. We would like to thank Eric Walther and Mareike Breda for assistance in recruiting and data collection. We would also like to thank Federico Meini and Tim Hollifield for their technical assistance. We thank all the children, and the staff from the Natural History Museum in Berlin for providing research space. The data is available on the Open Science Framework (OSF) at this link:

https://osf.io/e3mkc/?view_only=b77cf9a79fce437a8c57b97fffa04bb4.

Abstract

Changing one variable at a time while controlling others is a key aspect of scientific experimentation and is a central component of STEM curricula. However, children struggle to learn and implement this strategy. Why do children’s intuitions about how best to intervene on a causal system conflict with accepted scientific practices?

Interestingly, mathematical analyses have shown that controlling variables is *not always* the most efficient learning strategy, and that its effectiveness depends crucially on the “causal sparsity” of the problem, i.e. how many variables are likely to impact the outcome. We show that children as young as seven are sensitive to the causal sparsity of an unfamiliar causal system and use this information to tailor their testing strategies. We also show that they display some important sub-skills of CVS: planning and interpreting controlled experiments. Therefore, by middle childhood there should be a viable base upon which to build when teaching students CVS. Our analyses also help to clarify under what conditions controlling variables is actually a worthwhile approach to scientific inquiry as this is not always the case, a fact that might come as a surprise even to professional scientists.

Keywords: causal sparsity; causal learning; Interventions; scientific reasoning; CVS

Introduction

Imagine you are gifted some seeds for the very first time in your life: a little tomato plant! You want it to thrive, so you need to figure out what makes and keeps it healthy. How much sun, water and fertilizer does it need? This kind of task requires performing a series of unconfounded experiments to isolate and control the impact of the different variables under consideration (e.g., sun, water, and fertilizer) on the system (e.g., the health of the plant). For example, one might keep the amount of sun and water constant, modify the amount of fertilizer and see what happens. This approach is often referred to as the *Control of Variables Strategy* (CVS; Chen & Klahr, 1999; Klahr, Zimmerman, & Jirout, 2011; Kuhn & Brannock, 1977). Mastery of CVS is considered a hallmark of mature reasoning, and has become a crucial component of the Science, Technology, Engineering and Mathematics (STEM) curricula, featuring as one of the assessment criteria in national standards for science education (e.g., see National Academy National Academy of Sciences, 2013, p. 52). Mastering CVS requires mastering four sub-skills: *planning* a controlled experiment, *interpreting* such experiments, *identifying* them, and *understanding* their indeterminacy (Chen & Klahr, 1999; Schwichow, Christoph, Boone, & Härtig, 2016). Schwichow, Christoph, et al. (2016) determined that understanding that a confounded experiment results in invalid conclusions seems to be more difficult for 7- to 9th-graders than interpreting and identifying controlled experiments. They did not examine the difficulty of planning such an experiment relative to these other sub-skills, but earlier work by Bullock and Ziegler (1999) suggests that planning an experiment may also be more difficult for children than interpreting or identifying experiments.

Previous work on CVS has shown that children do not seem to be able to acquire CVS without instruction, even tend to manipulate multiple variables (Wilkening & Huber, 2004), and only start to be able to transfer it to new tasks from around age 10 (Chen & Klahr, 1999; Klahr, Fay, & Dunbar, 1993; Klahr et al., 2011; Kuhn, 2007; Kuhn

et al., 1995; Schauble, 1996; Wilkening & Huber, 2004; Zimmerman, 2007), despite the fact that they display some important precursor skills by age 6 or 7 in certain contexts (identifying and planning controlled tests; Sodian, Zaitchik, & Carey, 1991) and even preschoolers can use CVS as a domain-general strategy if given regular feedback and guidance (van der Graaf, Segers, & Verhoeven, 2015). This appears inconsistent with the growing number of studies investigating the developmental trajectory of causal learning, which show that toddlers' and preschoolers' active causal learning skills are, in some respects, already quite sophisticated (Cook, Goodman, & Schulz, 2011; Gopnik, Sobel, Schulz, & Glymour, 2001; Kushnir & Gopnik, 2005; Lucas, Bridgers, Griffiths, & Gopnik, 2014; McCormack, Bramley, Frosch, Patrick, & Lagnado, 2016; Ruggeri, Sim, & Xu, 2017; Schulz, Gopnik, & Glymour, 2007). Furthermore, the literature on CVS varies with respect to the type of instruction and specific tasks which afford students the best chance at learning this strategy, with some indications that direct instruction may be more beneficial, but which have not always been replicated (reviewed in Schwichow, Croker, Zimmerman, Höffler, & Härtig, 2016). It is therefore unclear from previous research whether children truly lack the necessary skills to spontaneously apply CVS or whether their documented difficulties in learning CVS may stem from something else. Evidence from studies on children's active learning and information search strategies, as well as the educational literature, suggests that children do indeed display some of the sub-skills that support learning and implementation of CVS, such as the abilities to plan, interpret, and possibly identify controlled experiments.

Children's effectiveness and adaptiveness in active learning

Causal understanding develops very early on in life. Children as young as two are already able to use patterns of variation and co-variation to infer causal relationships (Gopnik et al., 2001) and are more flexible than adults when inferring novel causal relationships (Lucas et al., 2014). Toddlers and preschoolers also spontaneously take on an active role in this type of exploration, making informative interventions to disambiguate the causal structure of a system, both in an experimental setting and

during spontaneous play (Cook et al., 2011; Kushnir & Gopnik, 2005; Schulz & Bonawitz, 2007; Sim & Xu, 2017), and the efficiency of these interventions increases with age (McCormack et al., 2016). More recently, developmental research has started investigating children’s effectiveness in active learning by looking at their ability to adapt their information search and hypothesis testing strategies to different characteristics of the system presented. This work showed that even 3- and 4-year-olds rely on different exploratory strategies depending on the statistical structure of the task at hand, *selecting* the strategy that maximizes information gain from among the given strategies (Ruggeri, Swaboda, Sim, & Gopnik, 2019). However, only by 7 years of age can children *generate* from scratch those actions that are most effective in a given situation, tailoring their spontaneous active learning strategies to the statistics of the given task (Ruggeri & Lombrozo, 2015).

Not all causal system are created equal: The role of causal sparsity

Several factors can (and should) impact children’s learning and hypothesis-testing strategies, such as the functional form of the causes under investigation, their relationship, and whether the causal learning system examined is deterministic or stochastic (see Schulz, Jones, Meder, and Ruggeri, 2020; Bonawitz, Denison, Gopnik, & Griffiths, 2014; Horn, Ruggeri, & Pachur, 2016; McCormack et al., 2016; Spiker & Cantor, 1979). In addition, it is important to take all information relevant to a problem into account when planning any inquiry strategy, as shown in previous work examining information search strategies. In this paper, we focus on one specific characteristic of the causal learning system, causal sparsity. Causal sparsity refers to the number of variables that are causally relevant for the system. As demonstrated by Coenen, Ruggeri, Bramley, and Gureckis (2019), causal sparsity mediates the effectiveness of different causal learning strategies, so that CVS is not always the most effective approach. In particular, Coenen et al. showed that as causal sparsity increases—that is, as the number of causes believed to be impacting a given outcome decreases—testing multiple variables at a time becomes more efficient. For example, imagine you want to

find out which of 20 switches on a poorly-labelled fusebox in the basement controls the bedroom light. Because only one switch controls the light, this system is very sparse. Implementing CVS, that is, testing one switch at a time, might take quite some time. To minimize the number of trips from the basement to the bedroom, it would make more sense to switch on half the switches on and see whether the bedroom light turns on (a so-called *Split half* strategy Nelson, Divjak, Gudmundsdottir, Martignon, & Meder, 2014). If it does, then the target switch must be one of those you had switched on; if it does not, it must be among the other set of switches. This way, you can rule out half of the remaining switches with each trip, instead of only one. However, even if you did not turn on exactly half of the remaining candidate switches for each trip, turning on multiple switches at a time is still faster than turning them on one by one as the number of trips up and down the stairs would still be reduced. Coenen et al. showed that adults were sensitive to the causal sparsity of a given system and adapted their strategies accordingly (although around $\frac{1}{3}$ defaulted to CVS irrespective of sparsity). This illustrates the importance of adaptiveness as a component of effective reasoning and learning. Yet to our knowledge, the adaptiveness of children's causal learning strategies has never been directly assessed in the classroom, where the focus has typically been on their mastery of CVS.

In this paper, we investigated to what extent 7- to 13-year-olds spontaneously display the ability to plan and interpret controlled experiments, and adapt their strategies to the context of a task to maximize testing efficiency. Based on previous work on children's information search skills, we expected children from age 7 to be sensitive to the statistical characteristics of a task, and be able to adapt their hypothesis-testing strategies accordingly, for instance by using CVS only when it is appropriate, depending on the causal sparsity of an unfamiliar system.

Method

Participants

Participants were 53 7- to 9-year-olds ($M = 8.19$ years, $SD = 0.59$, 24 female) and 51 10- to 13-year-olds ($M = 11.17$ years, $SD = 1.28$, 16 female), recruited and tested in museums in Berlin. All participants were German or fluent in German. IRB approval was granted and informed consent was obtained from parents prior to children's participation. All measures and conditions are described in this paper. The data is available on the Open Science Framework (Jones, Bramley, Gureckis, & Ruggeri, 2021): https://osf.io/e3mkc/?view_only=b77cf9a79fce437a8c57b97fffa04bb4.

Design and materials

We chose a paradigm which would allow us to assess children's ability to plan and interpret controlled experiments, as well as adapt their inquiries to different task contexts, in a task that could be completed hands-on and in a short time so as not to tax the youngest participants' attention spans, and which required minimal domain-specific knowledge, as it has been shown that causal learning tasks involving this type of knowledge are more challenging for children than those where only domain-general knowledge is required (Schwichow, Croker, et al., 2016; van der Graaf et al., 2015). Crucially, we needed a task that tested children's adaptiveness, allowing us to easily manipulate the efficiency of the strategies that could be implemented, so that CVS was not always the optimal approach. We therefore adapted the 'box-of-switches' task used by Coenen et al. (2019) in their work on adults' sensitivity to causal sparsity.

Participants were presented with a wooden box measuring approximately $35 \times 25 \times 10$ cm, the same used in Coenen et al. (2019). The top of the box featured six different toggle switches (corresponding to the six putative causes) on the left side, three lights (outcome), a red activation toggle and a slot to insert coin tokens (Figure 1). We limited the number of switches to 6 as the number of variables to be considered in a causal learning task is known to impact children's ability to use CVS successfully (Wilkening & Huber, 2004), and we wanted children to be able to complete

the task without assistance. Each switch could be turned on or off; depending on the experimental condition, different combinations of switches being on or off could turn on the lights. The activation toggle controlled whether the box was active or inactive. If it was inactive, the lights would never turn on. The box contained a raspberry Pi microcomputer (Richardson & Wallace, 2012) that determined the outcomes and recorded children's actions during the study.

Participants were randomly assigned to two conditions: *Sparse* and *Non-sparse*. In the Sparse condition, children were told that only one of the switches could turn on the lights. In the Non-sparse condition, they were told that all the switches could turn on the lights, except for one, which was broken. A single working switch was enough to turn on the lights. In both conditions, children's task was to find the unique (i.e., working or not working) switch. Which switch was working or not working was randomly determined for each child recorded to a tablet. In both conditions, the lights could only be turned on when the activation toggle was in the "on" position. Participants therefore set the switches in different positions, then turned on the activation toggle to see whether the outcome was present.

Procedure

Children were first familiarized with the box and its components. The experimenter explained the binary (left = off, right = on) nature of the switches and the difference between broken and working switches. Children were then instructed that they had to identify the working switch (in the Sparse condition) or the broken switch (in the Non-sparse condition), depending on the condition they were assigned to. Before starting the task, in both conditions, participants were led by the experimenter through two familiarization trials to practice the procedure and experience both outcomes: one in which they had to set all the switches to "on", and activate the box using the activation toggle, causing the light bulbs to turn on, and one in which they set all the switches to "off", activated the box again, and saw that the light bulbs did not turn on.

In subsequent trials, to identify the target switch, children could then test




Figure 1. Photograph of switch box used for the study.

different combinations of on/off switches and see if the light bulbs turned on as a result of each test. All switches were set to the “off” position again before the beginning of each new trial. To promote efficient search, participants were given six coin tokens at the beginning of the experiment, and had to pay one token using the slot provided (see Figure 1) every time they wanted to test a new switch combination. Participants could therefore perform up to six tests, but could stop at any time before then if they felt they had found the target switch. They were then asked to indicate to the experimenter which switch they thought was broken/working. The experimenter tested this by turning that switch on and activating the box so they could observe the outcome. If the child’s selection was correct, the experiment ended and they could keep their remaining tokens (each worth 0.50€). If not, they were given the option to perform more tests and guess again, or guess again right away, until they found the correct switch. The maximum reward was thus 2.50€, achievable if they were lucky enough to reach the solution after a single test trial. By following the ideal “split half” strategy it was possible to achieve $\approx 1.90\text{€}$ on average in the Sparse condition, while in the Non-sparse

condition, the only effective strategy was to test one switch at a time, with an expected return of 1.25€. If they used up all their tokens, or got the answer wrong, children received a sticker as a reward.

Results

Analysis of the first intervention

The number of switches tested in the very first intervention is crucially indicative of the way children approach the task in the different conditions (see Coenen et al., 2019). The number of children who tested one or multiple switches in each condition is shown in Table 1. We used logistic regression to evaluate to what extent Age group and Condition influenced the tendency to test one versus multiple switches in the first intervention. Both Age group (Odds Ratio [OR] = 2.52, [1.10, 5.78], $p = .03$) and Condition ($OR = 0.42$, [0.19, 0.97], $p = .04$) were significant predictors of whether children tested one switch at a time, with the relative odds of this being higher in younger children and in the Non-sparse condition. Including an Age group \times Condition interaction did not improve model fit ($\chi^2(1) = 0.003$, $p = .96$), indicating that age and condition contributed independently to children's tendency to test one switch in their first intervention. This suggests that children in both age groups were equally sensitive to causal sparsity, although their default strategy might shift with age, from testing one to testing multiple switches at a time.

Performance

Twelve children were excluded from subsequent analyses because their intervention data was incomplete due to technical difficulties, leaving 92 participants for whom we have a complete record. In total, 46 7- to 9-year-olds ($M = 8.21$ years, $SD = 0.55$, 20 female) and 46 10- to 13-year-olds ($M = 11.18$ years, $SD = 1.34$, 12 female) were included in the following analyses.

In the Sparse condition, 62% of younger participants (13/21) and 88% of older participants (23/26) identified the correct switch, having made a respective $M = 3.0$

Table 1

Counts and Percentage of Children Testing One or Multiple Switches on First Intervention.

Age group	Condition	Test One (first trial)	Test Multiple (first trial)
Younger	Sparse	15 (62.5%)	9 (37.5%)
	Non-sparse	23 (79.3%)	6 (20.7%)
Older	Sparse	11 (39.3%)	17 (60.7%)
	Non-sparse	14 (60.9%)	9 (39.1%)

($SD = 1.5$) and $M = 2.8$ ($SD = 1.3$) average interventions. In the Non-sparse condition, 64% of younger participants (16/25) and 50% of older participants (10/20) identified the correct switch, having made a respective $M = 3.5$ ($SD = 1.3$) and $M = 3.0$ ($SD = 1.5$) average interventions. χ^2 tests reveal that all age and condition combinations performed significantly above chance (all $ps < .001$).

Poisson regressions of the number of trials and the number of guesses children made with age group and condition as predictors showed no significant effects (see Appendix A1).

Analysis across all interventions

We also analyzed children’s sequences of interventions.

Strategy use. We classified children’s strategies into three types based on how many switches they turned on in each trial:

Test One denoted strategies in which exactly one switch was flipped on for every potentially informative intervention (see also Appendix A2). *Test Multiple* denoted strategies in which more than one switch was tested on every trial (see Appendix A2 for further details). Strategies which did not fit either of these criteria were classified as *Other*. The percentage of children who used each strategy, their accuracy, as well as strategy classifications from Coenen et al.’s adult sample, are shown in Figure 2. A more

detailed analysis of participants' strategy classification is presented in Appendix A2.

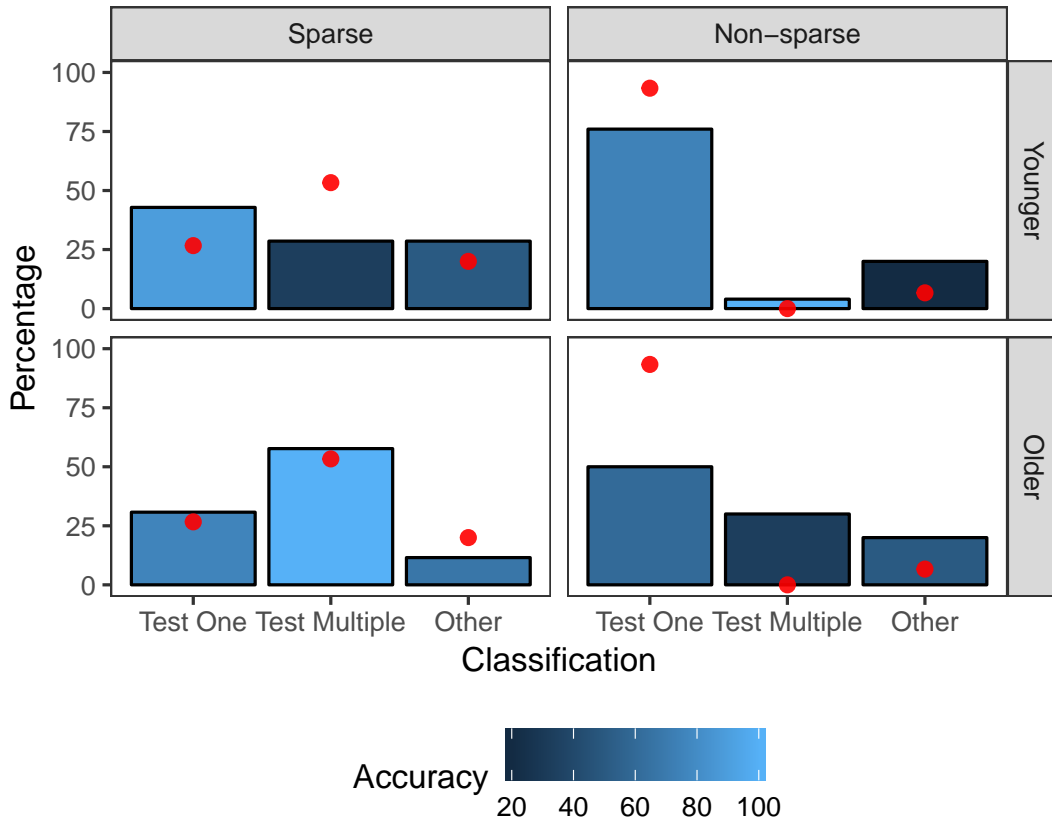


Figure 2. Bars show percentage of children in each age group classified as using each strategy in each condition. Shading shows the accuracy (percentage answering correctly) in each case. Red points represent data from Coenen et al.'s Experiment 1 adult sample ($N=30$) for comparison. Please note that the adults' strategy classifications also took selection EIG into account.

As for the analysis of the initial intervention, we used logistic regression to model whether Age group or Condition impacted on strategy classification. In line with the analyses of the first intervention, we found that older children were less likely to employ a Test One strategy, though this was only marginally significant ($OR = 0.44$, $[0.18, 1.04]$, $p = .06$), and that Test One was significantly more common in the Non-sparse condition ($OR = 3.06$, $[1.28, 7.30]$, $p = .01$). Including an Age group \times Condition interaction did not significantly improve model fit ($\chi(1) = -0.501$, $p = .48$). Older children were also significantly more likely to employ a Test Multiple strategy

($OR = 4.65$, $[1.64, 13.14]$, $p = .004$), and this strategy was significantly less common in the Non-sparse condition ($OR = 0.23$, $[0.08, 0.65]$, $p = .006$). Adding an Age group \times Condition interaction did not improve model fit ($\chi(1) = 0.81$, $p = .37$).

Strikingly, 17/21 (81%) of older children who classified as Test Multiple guessed the correct switch, while only 3/7 (42.9%) of younger participants classified as Test Multiple did so, though this difference did not reach significance (Fisher’s Exact Test, $p = .14$). In the Sparse condition, these proportions were 15/15 (100%) and 2/6 (33.3%), respectively (Fisher’s Exact Test, $p = .003$). Thus, together with our analysis of children’s first intervention, these results suggest that all children were sensitive to causal sparsity, although only older children were able to learn effectively from the more complex tests available in the Sparse condition. This is consistent with recent findings from Nussenbaum et al. (2020), who found that the ability to make causal inferences appears to develop separately from that to make appropriate interventions.

Expected information gain of children’s selections. The effectiveness of children’s interventions can also be explored using expected information gain (EIG). EIG is a common measure for how valuable information-seeking actions are to a learner, given their current state of uncertainty and learning goals (Nelson, 2005). A detailed explanation of how EIG is calculated can be found in Appendix A3. Here, the relative values of the available interventions are partly a function of learning condition. The Sparse condition has a wider range of actions that are potentially informative—any combination of between 1 and 5 switches is informative on the first test and many continue to be informative as the space of possibilities is narrowed, but within these options, choices that more evenly divide the remaining options are more informative than those that do so unevenly. In contrast, in the Non-sparse condition only a smaller range of interventions is informative—only those that turn on a single switch and have not already been performed.

To account for these differences, we computed the efficiency of each of participant’s interventions as a proportion of the most informative intervention available at that point from the perspective of an optimal learner that maximizes EIG at each step of the

search process, accurately integrating the evidence from all the previous interventions. As a baseline for comparison, we also simulated a set of learners that chose each intervention at random, flipping switches on with $p = .5$ but performing an equivalent total number of interventions as the participants. Figure 3 shows the efficiency of participants' interventions compared to those of the random baseline simulations. In all Age group and Condition combinations, interventions were significantly more informative than simulated random choices ($ps < .05$), with the exception of older children in the Sparse condition ($t(33) = 1.8, p = .069$). They were also significantly lower than the ceiling efficiency level of 1 (all $ps < .05$). Efficiency did not differ significantly by Age group or Condition, nor was there evidence for an interaction.

A more detailed analysis of children's strategy efficiency, that takes into account early stopping and unnecessary tests, is presented in A4.

Discussion

Designing a good experiment requires an understanding of the structure of the problem one wants to learn about. In this sense, no learning strategy is *always* best—not even the Control of Variables Strategy. In this study, we investigated whether and how 7- to 13-year-olds adapt their learning strategies and interventions to the characteristics of the causal learning system under investigation, focusing on its causal sparsity, and to what extent they can plan and learn from controlled experiments without guidance. Our results suggest that children may be already adaptive causal learners by age 7, and are almost as sensitive to causal sparsity as adults (Coenen et al., 2019). However, note that since children's strategy classifications did not take selection EIG into account, while the adults' did, it is difficult to directly compare the results from the two studies. This sensitivity echoes previous findings from the active learning literature, showing that children are ecological, adaptive learners from a very early age (Ruggeri & Lombrozo, 2015; Ruggeri et al., 2017, 2019). For example, Ruggeri and colleagues (2015) showed that in question-asking tasks, despite having difficulties generating effective questions from scratch until late childhood, children are as adaptive

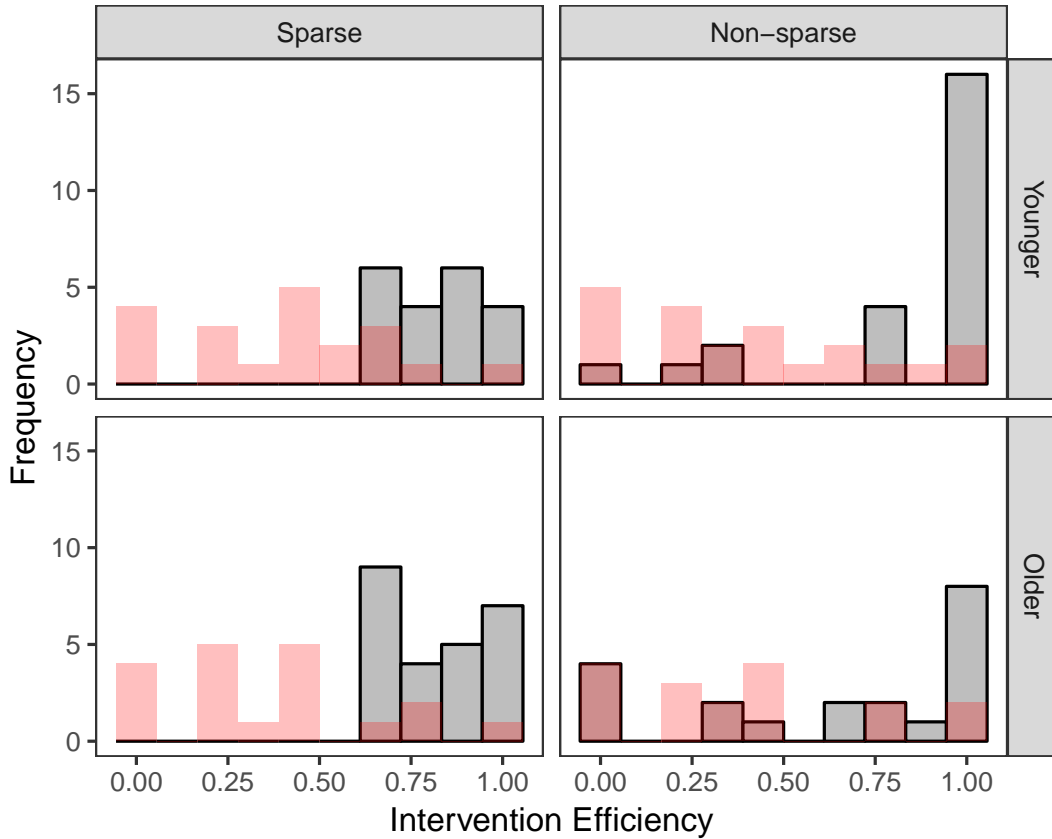


Figure 3. Grey bars indicate the efficiency of participants’ interventions, relative to an optimal intervener by Age group and Condition. Red bars show the efficiency of simulated learners that selected each intervention at random.

as adults (Ruggeri & Lombrozo, 2015).

In addition, we also found that children’s default strategy shifts with age from a Test One strategy to a Test Multiple strategy, with older children displaying patterns that more closely resembled those of adults (Coenen et al., 2019), but who nonetheless seemed to somewhat favor a Test One strategy. This is further supported by the prevalence of children whose strategies were classified as Other in the Non-sparse condition, as several of these children switched from testing one switch to multiple switches throughout the task (see Appendix A2), suggesting that they began with their default strategy but eventually realised that this was not an effective approach. This shows that children in this age range do seem to display the ability to plan and interpret controlled experiments in some situations, even without guidance.

Furthermore, this developmental difference in their default strategies appears to contradict previous evidence that young children tend to manipulate multiple variables at once even when they should not (Wilkening & Huber, 2004).

One other possible interpretation of these preferences could be that children's choices reflected an influence of *outcome valence*, as Tschirgi (1980) found that children and adults tend to test one variable at a time when the outcome is negative (e.g., when the light stays off), and change multiple variables when it is positive (e.g., when the light turns on), which roughly coincides with the most efficient strategies in the Non-sparse and Sparse condition, respectively. However, in Tschirgi's (1980) work, older children and adults showed a preference for testing one variable at a time, while younger children favored the opposite. Our participants displayed the opposite pattern of behavior, with younger children defaulting to a Test One strategy to a greater extent than older children. This suggests that outcome valence was not a major factor in children's strategy choices in our study.

This developmental trajectory is also strikingly similar to the development of children's question-asking strategies. In 20-Questions paradigms, children under the age of 7 almost exclusively ask *hypothesis-scanning* questions, which target one hypothesis at a time (e.g., "is it this one parrot?"). Between the ages of 7 and 10, children begin to ask more *constraint-seeking* questions, which target several hypotheses at once (e.g., "is it a bird?") by addressing categorical or perceptual features that are shared by several hypotheses, until this becomes the default strategy in adulthood (Herwig, 1982; Mosher & Hornsby, 1966; Ruggeri & Feufel, 2015). This suggests that children's learning strategies may broadly progress from being able to consider and reason about only one hypothesis at a time to being able to consider the entire range of hypotheses and their relationship with the outcome.

This may be due to a need to use strategies that respect the constraints of young children's cognitive abilities. Indeed, it is possible that in our study, younger children may have favoured CVS in part because a Test Multiple strategy was too resource-intensive, as keeping track of which switches were tested when, and updating

their probabilities accordingly, presumably requires greater executive function recruitment than CVS. Indeed, we know that children's ability to update multiple entries in working memory improves with age (Pailian, Carey, Halberda, & Pepperberg, 2020). Providing participants with memory aids to remind them of which switches they had tested previously may have helped lighten this burden and allowed us to test this hypothesis more directly. Along the same lines, a certain level of metacognitive skill is also very likely to play an important role in children's strategy choices, not only in tailoring their inquiries to their cognitive abilities, but also in explicitly selecting the best approach. Consistent with this idea, providing middle-schoolers with metacognitive training was found to benefit their science learning through improvements in their self-led learning processes (Zepeda, Elizabeth Richey, Ronevich, & Nokes-Malach, 2015). However, given the evidence that older children still sometimes defaulted to a Test One strategy, it seems unlikely that this was the only factor that influenced children's preferences.

It is an open question whether the default in hypothesis-testing strategies is driven by the prior assumptions children of different ages might have about the causal sparsity of a system (see Coenen et al., 2019). The fact that younger children performed slightly better than older children in the Non-Sparse condition may be reflective of a conflict between older children's prior beliefs and the instructions given about the causal learning system they were presented with. Consistent with this interpretation, some older children persisted with a Test Multiple strategy in this condition, although it was not effective due to the non-sparse nature of the problem.

Additionally, Coenen et al. (2019) showed that the effect of causal sparsity was strongly affected by the total number of variables in the system, that is, the more switches were presented to participants, the more prominent was participants' use of a Test Multiple strategy in the Sparse condition—as predicted by the model. In this sense, it would be interesting to test children on a version of this game with more switches. However, pilot testing indicated that this version of the task might be too challenging for children younger than 10.

Overall, our results are consistent with Sodian et al. (1991)'s findings, suggesting that children do display the ability to plan and interpret controlled experiments from the age of 7, though their ability to do so improves with age. Together with previous work which showed that young children can identify and interpret such experiments (Bullock & Ziegler, 1999), this indicates that three of the four basic sub-skills of CVS are present in some form by middle childhood, which should constitute a viable base for more explicit teaching of CVS. That it remains challenging for students to properly acquire CVS without extensive instruction is therefore puzzling.

However, it is possible that children's ability and propensity to use CVS in the appropriate context in our study was not an entirely deliberate choice but rather a reflection of their *default* approach. On the one hand, this interpretation does not fit with the work reviewed above, robustly showing how difficult it is for children to correctly implement and master CVS. On the other hand, children's use of CVS might reflect the fact that our task was relatively easy compared to the kinds of domain-specific CVS tasks that are taught in the classroom and usually investigated in the CVS literature. Our participants may have performed much worse when faced with such a task, which could have indicated that these sub-skills are not yet well-established at this age, or that domain-specific knowledge has a very great impact on strategy choices. Indeed, CVS and content knowledge are strongly interrelated (Edelsbrunner, Schalk, Schumacher, & Stern, 2018; van der Graaf et al., 2015). However, our task does not allow us to explore these possibilities.

Furthermore, our findings suggest that sensitivity to causal sparsity and strategy adaptiveness may also be worth considering when teaching STEM subjects, as CVS is not always the most effective approach. Children may sometimes believe that other strategies are more appropriate, which may account for some resistance to applying CVS to all the problems they are presented with. It may be more effective to provide children with a toolbox of strategies and teach them how and when to use each one, rather than focusing on training them to use only one strategy.

References

- Bonawitz, E., Denison, S., Gopnik, A., & Griffiths, T. L. (2014). Win-Stay, Lose-Sample: A simple sequential algorithm for approximating Bayesian inference. *Cognitive Psychology*. doi: 10.1016/j.cogpsych.2014.06.003
- Bullock, M., & Ziegler, A. (1999). Scientific reasoning: Developmental and individual differences. In F. E. Weinert & W. Schneider (Eds.), *Individual development from 3 to 12: Findings from the Munich longitudinal study* (pp. 38–54). New York, NY: Cambridge University Press.
- Chen, Z., & Klahr, D. (1999). All other things being equal: Acquisition and transfer of the control of variables strategy. *Child Development*, *70*(5), 1098–1120. doi: 10.1111/1467-8624.00081
- Coenen, A., Ruggeri, A., Bramley, N. R., & Gureckis, T. M. (2019). Testing one or multiple: How beliefs about sparsity affect causal experimentation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *45*(11), 1923–1941. doi: 10.1037/xlm0000680
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011, sep). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, *120*(3), 341–349. doi: 10.1016/J.COGNITION.2011.03.003
- Edelsbrunner, P. A., Schalk, L., Schumacher, R., & Stern, E. (2018). Variable control and conceptual change: A large-scale quantitative study in elementary school. *Learning and Individual Differences*, *66*(August), 38–53. doi: 10.1016/j.lindif.2018.02.003
- Gopnik, A., Sobel, D. M., Schulz, L. E., & Glymour, C. (2001). Causal learning mechanisms in very young children: two-, three-, and four-year-olds infer causal relationships from patterns of variation and covariation. *Developmental Psychology*, *37*(5), 620. doi: 10.1037//0012-1649.37.5.620
- Herwig, J. E. (1982). Effects of age, stimuli, and category recognition factors in children's inquiry behavior. *Journal of Experimental Child Psychology*, *33*(2), 196–206. doi: 10.1016/0022-0965(82)90015-7

- Horn, S. S., Ruggeri, A., & Pachur, T. (2016). The development of adaptive decision making: Recognition-based inference in children and adolescents. *Developmental Psychology, 52*(9), 1470–1485. doi: 10.1037/dev0000181
- Jones, A., Bramley, N. R., Gureckis, T. M., & Ruggeri, A. (2021). Changing many things at once sometimes makes for a good experiment, and children know that. *OSF Repository*. Retrieved from https://osf.io/e3mkc/?view_only=b77cf9a79fce437a8c57b97fffa04bb4 doi: 10.17605/OSF.IO/E3MKC
- Klahr, D., Fay, A. L., & Dunbar, K. (1993). Heuristics for scientific experimentation: A developmental study. *Cognitive Psychology, 25*(1), 111–146. doi: 10.1006/cogp.1993.1003
- Klahr, D., Zimmerman, C., & Jirout, J. (2011). Educational Interventions to Advance Children’s Scientific Thinking. *Science, 333*(6045), 971–975. doi: 10.1126/science.1204528
- Kuhn, D. (2007). Is direct instruction an answer to the right question? *Educational Psychologist, 42*(2), 109–113. doi: 10.1080/00461520701263376
- Kuhn, D., & Brannock, J. (1977). Development of the isolation of variables scheme in experimental and “natural experiment” contexts. *Developmental Psychology, 13*(1), 9. doi: 10.1037/0012-1649.13.1.9
- Kuhn, D., Garcia-Mila, M., Zohar, A., Andersen, C., White, S. H., Klahr, D., & Carver, S. M. (1995). Strategies of knowledge acquisition. *Monographs of the Society for Research in Child Development, 60*(4), 1–157. doi: 10.2307/1166059
- Kushnir, T., & Gopnik, A. (2005). Young children infer causal strength from probabilities and interventions. *Psychological Science, 16*(9), 678–683. doi: 10.1111/j.1467-9280.2005.01595.x
- Lucas, C. G., Bridgers, S., Griffiths, T. L., & Gopnik, A. (2014). When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition, 131*(2), 284–299. doi: 10.1016/j.cognition.2013.12.010

McCormack, T., Bramley, N. R., Frosch, C., Patrick, F., & Lagnado, D. A. (2016).

Children's use of interventions to learn causal structure. *Journal of Experimental Child Psychology*, *141*, 1–22. doi: 10.1016/j.jecp.2015.06.017

Mosher, F. A., & Hornsby, J. R. (1966). On asking questions. *Studies in Cognitive Growth*, 86–102.

National Academy of Sciences. (2013). *Next generation science standards: For states, by states*. Washington, DC: National Academies Press.

Nelson, J. D. (2005). Finding useful questions: On bayesian diagnosticity, probability, impact, and information gain. *Psychological Review*, *112*(4), 979–999. doi: 10.1037/0033-295X.112.4.979

Nelson, J. D., Divjak, B., Gudmundsdottir, G., Martignon, L. F., & Meder, B. (2014). Children's sequential information search is sensitive to environmental probabilities. *Cognition*, *130*(1), 74–80. doi: 10.1016/j.cognition.2013.09.007

Nussenbaum, K., Cohen, A. O., Davis, Z. J., Halpern, D. J., Gureckis, T. M., & Hartley, C. A. (2020). Causal Information-Seeking Strategies Change Across Childhood and Adolescence. *Cognitive Science*, *44*(9). doi: 10.1111/cogs.12888

Pailian, H., Carey, S. E., Halberda, J., & Pepperberg, I. M. (2020). Age and Species Comparisons of Visual Mental Manipulation Ability as Evidence for its Development and Evolution. *Nature: Scientific Reports*, *10*, 7689. doi: 10.1038/s41598-020-64666-1

Richardson, M., & Wallace, S. (2012). *Getting started with raspberry pi*. O'Reilly Media, Inc.

Ruggeri, A., & Feufel, M. (2015). How basic-level objects facilitate question-asking in a categorization task. *Frontiers in Psychology*, *6*, 918. doi: 10.3389/fpsyg.2015.00918

Ruggeri, A., & Lombrozo, T. (2015). Children adapt their questions to achieve efficient search. *Cognition*, *143*, 203–216. doi: 10.1016/j.cognition.2015.07.004

Ruggeri, A., Sim, Z. L., & Xu, F. (2017). “Why Is Toma Late to School Again?”: Preschoolers Identify the Most Informative Questions. *Developmental Psychology*,

- 53(9), 1620–1632. doi: 10.1037/dev0000340
- Ruggeri, A., Swaboda, N., Sim, Z. L., & Gopnik, A. (2019). Shake it baby, but only when needed: Preschoolers adapt their exploratory strategies to the information structure of the task. *Cognition*, 193, 104013. doi: 10.1016/j.cognition.2019.104013
- Schauble, L. (1996). The development of scientific reasoning in knowledge-rich contexts. *Developmental Psychology*, 32(1), 102–119. doi: 10.1037/0012-1649.32.1.102
- Schulz, L. E., & Bonawitz, E. B. (2007). Serious Fun: Preschoolers Engage in More Exploratory Play When Evidence Is Confounded. *Developmental Psychology*, 43(4), 1045–1050. doi: 10.1037/0012-1649.43.4.1045
- Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science*, 10(3), 322–332. doi: 10.1111/j.1467-7687.2007.00587.x
- Schwichow, M., Christoph, S., Boone, W. J., & Härtig, H. (2016). The impact of sub-skills and item content on students' skills with regard to the control-of-variables strategy. *International Journal of Science Education*, 38(2), 216–237. doi: 10.1080/09500693.2015.1137651
- Schwichow, M., Croker, S., Zimmerman, C., Höffler, T., & Härtig, H. (2016, mar). Teaching the control-of-variables strategy: A meta-analysis. *Developmental Review*, 39, 37–63. doi: 10.1016/j.dr.2015.12.001
- Shannon, C. E. (1951). Prediction and entropy of printed english. *The Bell System Technical Journal*, 30, 50–64.
- Sim, Z. L., & Xu, F. (2017). Learning higher-order generalizations through free play: Evidence from 2- and 3-year-old children. *Developmental Psychology*, 53(4), 642–651. doi: 10.1037/dev0000278
- Sodian, B., Zaitchik, D., & Carey, S. (1991). Young Children's Differentiation of Hypothetical Beliefs from Evidence. *Child Development*, 62(4), 753–766. doi: 10.1111/j.1467-8624.1991.tb01567.x
- Spiker, C. C., & Cantor, J. H. (1979, oct). Factors affecting hypothesis testing in

- kindergarten children. *Journal of Experimental Child Psychology*, *28*(2), 230–248.
doi: 10.1016/0022-0965(79)90086-9
- Tschirgi, J. E. (1980). Sensible Reasoning: A Hypothesis about Hypotheses. *Child Development*, *51*(1), 1–10. doi: 10.2307/1129583
- van der Graaf, J., Segers, E., & Verhoeven, L. (2015). Scientific reasoning abilities in kindergarten: dynamic assessment of the control of variables strategy. *Instructional Science*, *43*(3), 381–400. doi: 10.1007/s11251-015-9344-y
- Wilkening, F., & Huber, S. (2004). Children's intuitive physics. In U. Goswami (Ed.), *The blackwell handbook of childhood cognitive development*. Blackwell Reference Online. doi: 10.1111/b.9780631218418.2004.00019.x.
- Zepeda, C. D., Elizabeth Richey, J., Ronevich, P., & Nokes-Malach, T. J. (2015). Direct instruction of metacognition benefits adolescent science learning, transfer, and motivation: An in vivo study. *Journal of Educational Psychology*, *107*(4), 954–970. doi: 10.1037/edu0000022
- Zimmerman, C. (2007, jun). The development of scientific thinking skills in elementary and middle school. *Developmental Review*, *27*(2), 172–223. doi: 10.1016/J.DR.2006.12.001

Appendix

A1. Number of interventions and guesses by Age Group and Condition

Table A1

Poisson Regressions predicting Number Interventions and Guesses by Age Group and Condition

	N Interventions			N Guesses		
	Mean	95% CI	<i>p</i>	Mean	95% CI	<i>p</i>
Intercept	3.00	2.34–3.84	< .001	1.33	0.92–1.93	.13
Age group (Younger)	0.94	0.67–1.31	.70	0.81	0.48–1.36	.42
Condition (Sparse)	1.17	0.85–1.62	.33	1.08	0.66–1.77	.76
Age group × Condition	0.91	0.57–1.46	.70	1.46	0.72–2.95	.29

Note: Coefficients and confidence intervals transformed to natural odds ratios.

Reference groups for factors indicated in brackets

A2. Strategy classification details

Coenen et al. (2019) explored adult switch box behaviour in cases with up to 20 switches. However, these strategy classifications also took the EIG of each selection into account. This allowed them to distinguish between 6 strategy classifications: *Pure Test One*, *Noisy Test One*, *Pure Test Half*, *Pure Test Multiple* and *Noisy Test Multiple*, with *Other* as a catch all. However for the 6 switch case explored here and in (Coenen et al., 2019, *Experiment 1*) *Test Half* and *Test Multiple* are not clearly distinguished.

Therefore following Coenen et al. (2019) *Experiment 1*, though we did not account for EIG when classifying children's strategies, we collapsed the 4 possible strategies in our study (*Test One*, *Test Half*, *Test Multiple* and *Other*) into 3 broader strategy categories (incorporating *Test Half* into *Test Multiple*). Below we provide the precise criteria for each fine grained category and show the full breakdown in Figure A1.

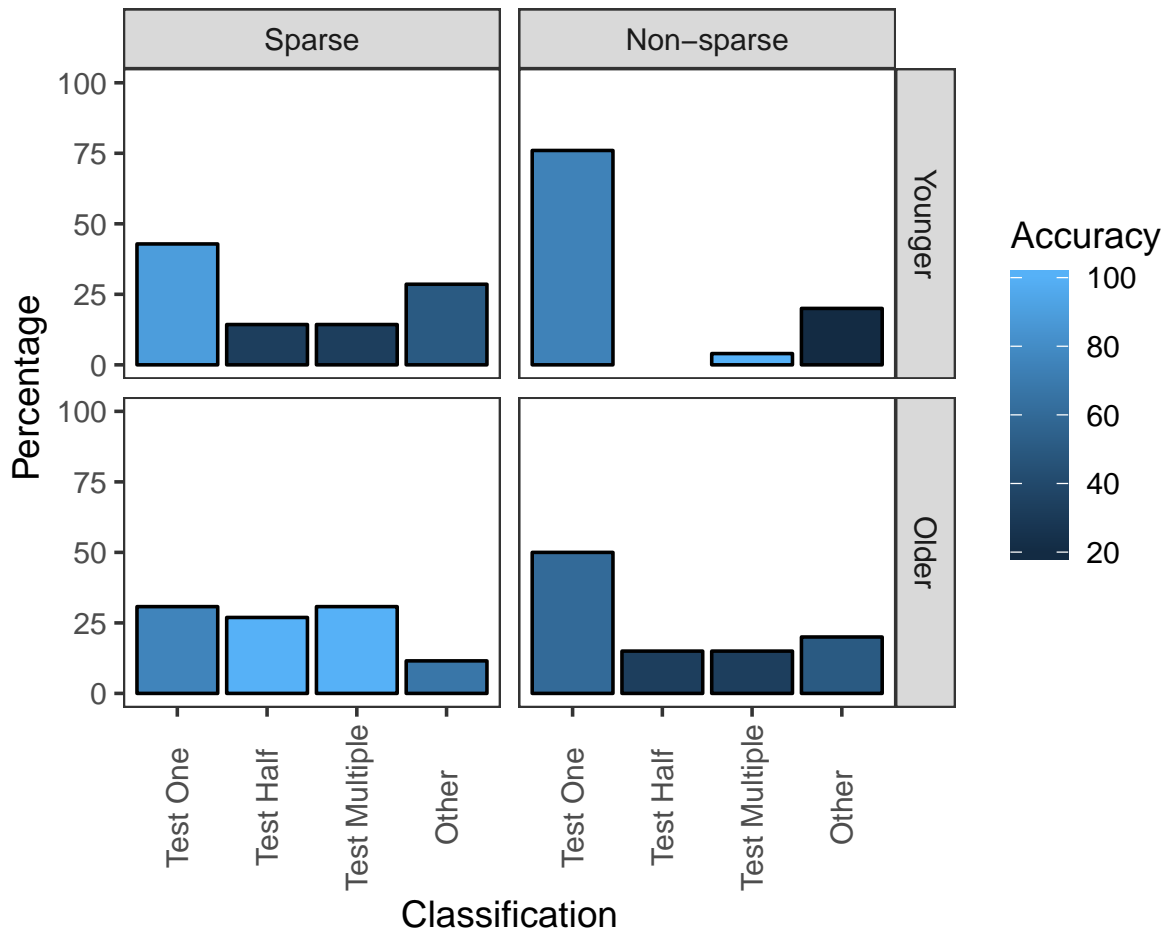


Figure A1. Fine grained strategy classifications by Age group and Condition. Shading shows the accuracy (percentage answering correctly) for each coarse grained strategy class, as in Figure 2.

1. **Test One:** The participant always turned one switch on at a time.
2. **Test Multiple:** The participant turned on several (but not always exactly half) of the remaining potential causes on every trial. *Test Half:* The participant manipulated exactly half of the remaining potential causes on every trial (rounding odd numbers up or down).
3. **Other:** Any strategy that does not fall into the above categories. This also included participants who switched back and forth between Testing One and Testing Multiple. For example, this category would include participants who started testing variables one-by-one and then changed their strategy to changing

half of the variables, and vice versa.

An ideal information-gain-maximising learner would always be classified as following a *Test Half* or *Test Multiple* strategy in the Sparse condition and a *Test One* strategy in the Non-sparse condition.

The Other classification included children who switched back and forth between a Test One and a Test Multiple strategy, but also children who started with a Test Multiple strategy before switching to Test One, or vice-versa, and children whose last test manipulated only one switch (e.g., for a confirmatory test) but who otherwise followed a Test Multiple approach. A total of 18 children were classified as Other; 9/18 (50%) in the Sparse condition, and 8/18 (44.4%) in the Non-sparse condition. Of these, 11 were younger and 7 were older. Ten (7 younger and 3 older) of the children classified as Other made selections consistent with beginning with one strategy (e.g., Test One or Test Multiple) but then switching to another, or manipulating a single switch on the last test while having followed a Test Multiple strategy until then. Most of these cases were found in the Non-sparse condition (6/10 or 60%). Only younger children were classified as this ‘sub-type’ of Other in the Sparse condition, and the proportion of children who were classified this way in the Non-sparse condition was the same across age groups (3 older children, 3 younger children). This may reflect both the relatively more difficult nature of the task in the Non-sparse condition, and perhaps a certain level of sensitivity to causal sparsity in younger children, most of whom may not have consistently tested several switches in the Sparse condition, but several of whom might have adopted this approach after realising during the task that it was more efficient.

A3. Expected information gain calculation

In this task, the learner is confronted with a causal system with $N = 6$ binary independent variables, I , of which a subset of variables $C \subseteq I$ (i.e., individual switches) can affect the outcome when active (i.e., switched to the “on” position, and one binary outcome, o (i.e., the lights turning on). The probability of the outcome given a specific setting of variables is

$$P(o = 1|C) = \begin{cases} 1, & \text{if } \exists c \in C \wedge (c = 1), \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Simply put, the outcome occurs if, and only if, any of the variables in C are currently active. The learner must decide how to manipulate the variables to determine which are causally relevant. We assume that the learner's optimal strategy consists of choosing a switch setting, $s \in S$, which maximizes the *Expected Information Gain* (EIG) with respect to the system. EIG quantifies the expected reduction in uncertainty over the hypotheses H after having made an intervention on the system and observed an outcome. Here, the learner's hypotheses are possible sets of causally relevant variables, i.e., $H = \{C_1, \dots, C_6\}$. Note that the contents of H differ between conditions because of the differences in sparsity. In the Sparse condition, each set (e.g., C_1) contains only one switch because only one switch can activate the lights, while in the Non-sparse condition, each C contains a combination of 5 switches, as all but one switch can turn on the lights. We consider a simple case of binary outcomes ($o = 1$ or $o = 0$) with the likelihood of an outcome given by Equation 1. A learner's EIG is calculated as

$$\text{EIG}(s|H) = \text{SE}(H) - \sum_{j=0}^1 P(o = j|s) \text{SE}(H|s, o = j), \quad (2)$$

where SE represents the Shannon entropy over a distribution of hypotheses (Shannon, 1951), which in this study are possible causes of the light turning on. The marginal likelihood of each outcome is then given by

$$P(o = j|s) = \sum_{i=1}^6 P(o = j|C_i; s) \quad (3)$$

and the prior entropy (i.e., the uncertainty as to whether each candidate hypothesis is correct before a test) is

$$\text{SE}(H) = - \sum_{i=1}^6 P(C_i) \log P(C_i). \quad (4)$$

After observing the outcome of a test, the learner's beliefs about each hypothesis are updated following Bayes' rule,

$$P(C_i|o) = \frac{P(o|C_i)P(C_i)}{\sum_{j=1}^6 P(o|C_j)P(C_j)}, \quad (5)$$

and the entropy over the updated set of hypotheses becomes

$$SE(H|s, o) = - \sum_{i=1}^6 P(C_i|o) \log P(C_i|o). \quad (6)$$

A4. Early stopping and unnecessary tests

Stopping early and making unnecessary tests are two kinds of search errors that can provide additional insight into the quality of a learner’s search. Stopping one’s search before identifying the correct switch may occur if a participant searched inefficiently and runs low on tests and chooses to guess. However guessing before not have an explicit understanding of the task or strategy, or that they might be using a more heuristic approach rather than consciously following a specific search strategy. Making “unnecessary tests”, that is, tests that occur after the target switch could have been identified and which therefore did not provide any additional information from a normative perspective, suggests that children may find it difficult to keep track of the evidence gathered previously, or that they don’t believe a trial fully rules out a switch setting.

The number and percentages of children stopping early and performing unnecessary tests is shown in Table A2. Performing unnecessary tests was rare, with only three children performing (either one or two) unnecessary tests. However, stopping early was common in both conditions for younger children and just in the Non-sparse condition for older children.

A logistic regression predicting early stopping with Age Group and Condition showed that the odds of stopping early did not differ significantly between conditions $OR = 1.02, [0.32, 3.24], p = .97$ but older children were less likely to stop early ($OR = 0.14, [0.03, 0.63], p = .001$), and that there was a significant interaction between Age Group and Condition, with older children much more likely to stop early in the Non-sparse condition ($OR = 14.02, [2.08, 94.5], p = .007$, see Table A2). For comparison, our random intervention baseline simulations produced test sequences that rarely resolved all uncertainty by the time participant made their judgment, effectively being classified as stopping early 60% of the time in the Sparse condition, and 89% of the

Table A2

Counts and Percentage of Children Stopping Early and Number of Unnecessary Tests Performed, and Average Number of Total Tests Performed (SD).

Age group	Condition	N Participants	Stopped Early	Tested Unnecessarily	N tests
Younger	Sparse	21	10 (47%)	1 (5%)	3.00 ± 1.94
	Non-sparse	25	12 (48%)	1 (4%)	3.52 ± 1.33
Older	Sparse	26	3 (12%)	1 (4%)	2.81 ± 1.30
	Non-sparse	20	13 (59%)	0 (0%)	3.00 ± 1.52

time in the Non-sparse condition.