

Learning Causal Overhypotheses through Exploration in Children and Computational Models

Jiakun Liu



Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2022-134

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2022/EECS-2022-134.html>

May 17, 2022

Copyright © 2022, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

I would first like to thank my research advisor Professor John Canny and my mentor David Chan. I would also like to thank my collaborators Eliza Kosoy, Jasmine Collins, Bryanna Kaufmann, Rosemary Nan Ke, Jessica B Hamrick and Sandy Han Huang for their contributions, and professor Alison Gopnik for her guidance. Finally, I would also like to thank research assistants Andy Liu, Janie Dent, Athena Leong, Eli Phipps, Jenna Levin, Nikita Kumar and Zane Levine for their meticulous work

Learning Causal Overhypotheses through Exploration in Children and Computational Models

by Adrian Liu

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences,
University of California at Berkeley, in partial satisfaction of the requirements for the
degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination:

Committee:

Professor John Canny
Research Advisor



(5/17/22)

Professor Alison Gopnik
Second Reader



May 11, 2022

Learning Causal Overhypotheses through Exploration in Children and Computational
Models

by

Adrian Liu

A technical report submitted in partial satisfaction of the

requirements for the degree of

Master of Science

in

Electrical Engineering and Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor John Canny, Chair

Professor Alison Gopnik

Spring 2022

Learning Causal Overhypotheses through Exploration in Children and Computational
Models

Copyright 2022
by
Adrian Liu

Abstract

Learning Causal Overhypotheses through Exploration in Children and Computational Models

by

Adrian Liu

Master of Science in Electrical Engineering and Computer Science

University of California, Berkeley

Professor John Canny, Chair

Despite recent progress in reinforcement learning (RL), RL algorithms for exploration still remain an active area of research. Existing methods often focus on state-based metrics, which do not consider the underlying causal structures of the environment, and while recent research has begun to explore RL environments for causal learning, these environments primarily leverage causal information through causal inference or induction rather than exploration. In contrast, human children—some of the most proficient explorers—have been shown to use causal information to great benefit. In this work, we introduce a novel RL environment designed with a controllable causal structure, which allows us to evaluate exploration strategies used by both agents and children in a unified environment. In addition, through experimentation on both computation models and children, we demonstrate that there are significant differences between information-gain optimal RL exploration in causal environments and the exploration of children in the same environments. We conclude with a discussion of how these findings may inspire new directions of research into efficient exploration and disambiguation of causal structures for RL algorithms.

Contents

Contents	i
List of Figures	ii
List of Tables	iv
1 Introduction	1
2 Related Work	4
3 The Virtual “Blicket Detector” Environment	6
3.1 Measuring Childrens’ Causal Exploration	7
4 Modeling Causal Learning Using Blickets	14
5 Discussion & Conclusion	20
Bibliography	22

List of Figures

3.1	Screenshot of our virtual blicket detector. Children can interact with the blicket detector through a touch interface on an iPad, or through a point and click interface on any web-enabled machine.	6
3.2	Visualization of the demonstration the children see for which objects make the blicket detector light up, per condition, either Condition 1: GIVEN HYPOTHESIS or Condition 2: NOT GIVEN HYPOTHESIS.	9
3.3	Left: Children’s verbal replies about objects’ blicket-ness. Blue bars indicate the true positive rate, i.e. the proportion of times children said that an object was a blicket, given that it was a blicket. Yellow bars indicate the false positive rate, i.e. the proportion of times children said that an object was a blicket, given that it was not. The error bars indicate standard error. Random guesses would achieve 50% for both true positive rate and false positive rate. The differences between all the conditions are statistically significant with $p < 0.02$, except between the CONJUNCTIVE GIVEN HYPOTHESIS conditions. Right: The percentage of children who generated enough data to make a valid conclusion about which objects are blickets (assuming an optimal inference procedure).	11
3.4	The average number of actions taken per condition, and the time that children require to make the first successful action (lighting up the detector) per causal model and condition.	11
4.1	The optimal policy for the PER-STEP model under UNIFORM prior. The observation tree is represented from left to right, with the first column representing the first action taken and the second column representing the resulting observation outcomes and subsequent action taken by the PER-STEP model. The model tests single objects, and if both tested objects are off, tries combinations of objects to attempt to reduce the uncertainty. The model is done when it knows the exact condition of the space.	16
4.2	The optimal policy for the PER-STEP model under the EXPERIMENTAL prior. This choice of prior reduces the number of policy branches, but also increases the average time to solution.	16
4.3	The optimal policy for the MINIMUM-STEP model under the UNIFORM prior. . .	16
4.4	The optimal policy for the MINIMUM-STEP model under EXPERIMENTAL prior.	17

- 4.5 Left: The average number of combinations the children tried before they indicated they finished exploring, compared to the number of combinations the models needed to fully determine the ground truth. Agents use the EXPERIMENTAL prior. We find that in the CONJUNCTIVE GIVEN HYPOTHESIS condition the children’s average number of combinations is statistically significant compared to the MINIMUM-STEP model, but the difference between the PER-STEP model is not significant. In the DISJUNCTIVE GIVEN HYPOTHESIS condition the difference between the children and the MINIMUM-STEP and PER-STEP models is significant with $p < 0.0001$. Right: The percentage of children who generated enough observations to completely determine the ground truth in each condition. While conj. vs. disj. have no significant difference ($p > 0.9$), when given the hypothesis space, children are more likely to disambiguate the causal space ($p = 0.002$). The models always generate enough observations to disambiguate the causal space. 18
- 4.6 Example sequences of what objects children placed and tested on the blicket detector. The 0 column gives the ground truth blicket objects. Each row represents the order the blickets are placed on the detector, and the final row is green if the blicket detector is illuminated. The gray column represents when the child has finished exploring and is asked which objects are blickets. 19

List of Tables

3.1	Statistics of children's exploration. Shown are the average number of checks (i.e., presses of the check mark) taken per condition, the average number of unique combinations tried per condition, the average time played per condition in seconds, and the average time played before seeing the blicket detector go on for the first time. Standard deviations are given in parenthesis.	9
3.2	Percentage of children who tested for hypotheses that should be ruled out in the GIVEN HYPOTHESIS conditions. If the children acted completely rationally, the first column should be 0 for all 4 hypotheses. Instead, we see the children are testing these hypotheses with near identical frequency regardless of the demonstration video they received	13

Acknowledgments

I would first like to thank my research advisor Professor John Canny and my mentor David Chan. I would also like thank my collaborators Eliza Kosoy, Jasmine Collins, Bryanna Kaufmann , Rosemary Nan Ke, Jessica B Hamrick and Sandy Han Huang for their contributions, and professor Alison Gopnik for her guidance. Finally, I would also like to thank research assistants Andy Liu, Janie Dent, Athena Leong, Eli Phipps, Jenna Levin, Nikita Kumar and Zane Levine for their meticulous work.

Chapter 1

Introduction

Exploration is a fundamental problem in reinforcement learning (RL). In order to act on the world effectively, an agent needs to be able to efficiently and actively gather information about how the environment works. Gathering causal information is particularly helpful for action planning and generalization [22, 44, 24]. For example, if the agent’s task is to turn on a lamp, the corresponding causal relationship is that the lamp will turn on only if all of the following are true: 1) the switch is flipped to the on position, 2) the lamp is plugged into a source of electricity, and 3) the light bulb is working. Understanding this causal relationship enables the agent to systematically diagnose a problem—if, for instance, it is in a new room and flips the switch to on but the lamp does not turn on—and systematically explore to find solutions.

Existing RL exploration methods typically do not focus on such *causal* exploration: they do not form and test causal hypotheses, or plan active interventions to obtain causal data [3]. Instead, existing RL exploration methods primarily focus on expanding the set of experiences of the agent, for instance by visiting novel or surprising areas of the state space. This may be sufficient for the agent to solve the particular task it is trained for, but limits its ability to generalize to new tasks and environments [41, 10]. Although there is growing interest in causal learning in RL, this work has focused on extracting a particular underlying causal graph from given data in a particular environment [35, 24, 50]. There is very limited amount work that attempts to utilize causal information for exploration in RL, or to learn abstract causal structure through exploration.

How might we integrate causal learning and reasoning into exploration in RL agents? We propose to draw inspiration from cognitive science. In contrast to RL agents, even young children learn and reason about causal relations and actively explore to collect causal data. Moreover, they can learn and use *causal overhypotheses*—hypotheses about which classes of causal relationships are more or less likely [25, 28]. Causal overhypotheses are a key component that allows humans to learn causal models from a sparse amount of data [21], because they can help narrow down the possible causal relationships that we consider and test.

In the previous light example, suppose that we enter a new apartment, and want to turn

on the lights in the living room. If we lacked any causal overhypothesis about what causes a light to turn on, then this would be an essentially hopeless task. We might try knocking on the wall, shuffling on the carpet, rotating the lamp, and so forth—the way that an RL agent starts out acting in an environment in which it has no prior experience. Instead, if our causal overhypothesis is that flipping a light switch on the wall will turn on one or more lights, then that helps our hypothesis testing. We might try flipping various combinations of light switches in the apartment, and quickly identify the correct causal relationships.

In our experiments, we draw inspiration from the blicket detector experiment [16], which is a classic setting for evaluating causal learning and reasoning in children. Blocks are placed on a “blicket machine”. Some blocks are “blickets” and the blicket machine lights up when blickets are placed on it, according to some rule. The participants must learn which blocks are blickets, and use those blocks to activate the machine.

The blicket machine requires children to learn the structure of a novel causal system, and allows researchers to present children with relatively complex patterns of statistical correlation and intervention. It does so in a concrete, simple and intuitive way. As a result, a large body of studies using this method have demonstrated a remarkable range of causal inference capacities in children as young as 18 months [19, 12, 14, 18, 15, 17, 26, 28, 33]. In particular, studies have tested how well children learn more abstract overhypotheses about the rules by which the machine works. For example, the causal relationship can be either disjunctive or conjunctive [28]. In the disjunctive case, if at least one blicket is placed on the machine, then it lights up. In the conjunctive case, at least two blickets must be placed on the machine in order for it to light up. Participants must infer whether the machine is conjunctive or disjunctive. Interestingly, prior work has found not only that children can learn these overhypotheses from data, but also that they are more flexible than adults in these tasks. They are better at learning unusual overhypotheses, like those involving conjunctive causal relationships [28, 20]

However, in prior experiments investigating causal overhypotheses, children are given the relevant data by the experimenter rather than generating the data themselves through causal exploration. Could children also actively generate data that would allow them to learn the right causal overhypotheses? And would causal overhypotheses shape their exploration? Developmental psychology has found that children are active and curious learners, with strong intrinsic motivation to systematically explore their environment [47, 46]. Even young children engage in hypothesis-testing behavior in settings with ambiguous [12] or inconsistent [27] evidence. However, there is limited prior work on spontaneous causal exploration in children and none on how causal overhypotheses affect children’s exploration.

In this work, we seek to understand how young children, between the ages of four and six, explore and learn causal structure and how their exploration compares to that exhibited by typical ideal observer or RL models. In particular, we investigate how children use exploration to learn causal overhypotheses and how causal overhypotheses influence their exploration. We first show them specific sequences of interactions with the blicket machine, which are consistent with either a disjunctive, conjunctive, or ambiguous causal relationship. We then give them as much time as they want to experiment with a new blicket machine, to

“figure out how to make it go.” We find that children exhibit rich and diverse exploratory behavior, and that their overhypotheses are indeed influenced by their prior exposure to the environment. In contrast, we consider an ideal learner with a particular causal overhypothesis and find that its exploratory behavior in this task is quite different from that of children. We conclude by discussing how our findings can inform the development of RL agents that are capable of causal exploration, learning and reasoning. Our findings cannot be directly applied to existing RL algorithms, but rather may inspire entirely new directions of research.

To summarize, the main contributions of this report are as follows. 1) We develop an online environment based on the classical blicket machine [16] which can be used with both children and agents. 2) We collect experimental data from children in this environment, and demonstrate that they exhibit diverse structured exploration strategies, and that they learn from this exploration. 3) We compare children’s exploration strategies with a set of simple ideal observer models, and show that children’s actions do not directly reflect either simple overhypothesis information gain or reward maximization. The results suggest that children rely on a broad set of causal assumptions and exploration behaviors that may generalize to many environments. This paves the way for future research which will enable artificial agents to exhibit richer, causally motivated, exploration strategies.

Chapter 2

Related Work

The relevant previous work includes studies of exploration in reinforcement learning (RL), multi-task RL, causal learning in RL, and various versions of the blicket environment used in cognitive science experiments.

Exploration in reinforcement learning Causal exploration is a relatively understudied area in reinforcement learning. Most techniques do not consider explicit causal hypotheses, but instead rely on adding an exploration bonus to the task reward. This exploration bonus may be given for visiting novel states [5, 40, 31, 49, 29], surprising dynamics [45, 43], uncertainty [39, 7] or disagreement [42]. Please refer to [3] for a comprehensive survey of exploration in deep reinforcement learning. Our proposed work using the blicket environment lays the groundwork for potential exploration algorithms based on the empirical exploration patterns of children in causal environments.

Multi-task learning in reinforcement learning There are several benchmarks for multi-task learning for robotics [51, 23], for physical reasoning [2, 4] and video games [9, 30, 37, 8]. The relevant causal overhypotheses for these environments are not clear, however, making it difficult to evaluate the influence of causal information on agents' exploration. In our work, we introduce a novel blicket environment, where the causal overhypotheses are clearly defined.

Causal learning in reinforcement learning In recent years several reinforcement learning benchmarks and environments for causal discovery have been proposed. The first group provide realistic simulations that capture of causal relationships in the real world. Causal World [1] provides simulations where the agent controls a robotic arm to perform a distribution of tasks such as push, picking and stacking blocks. It allows the experimenter to design a wide range of tasks that share the same causal structure. Similarly, Causal City [32] provides a realistic navigation environment for driving simulation. It contains rich causal information in the form of behaviors of drivers, pedestrians, traffic lights and other environment features.

Both of these examples involve different tasks with shared causal structures that are fixed, such as laws of physics or the behavior of pedestrians, and are primarily concerned with how causal knowledge learned from one task can be transferred to another. This encourages the agent to store specific causal knowledge in its network parameters, rather than learn an exploration strategy so that it can discover the causal structure in an unseen task.

Another group of benchmarks and environments are more closely related to our work in that they use a meta-RL paradigm where the causal structure is randomly generated for each episode. Thus the agent needs to determine the causal structure for every task it sees. *Alchemy* [50] involves a task where the agent plays the role of a chemist, and needs to figure out what each potion does to each stone. The specific set of rules for which potions interact with stones is called a "chemistry", and is randomly generated each episode. More interestingly, *Alchemy* provides two different interfaces for the agent: an embodied version where the agent needs to visually understand the environment and navigate through space to interact with the objects it wants, and a symbolic version where both the observation space and action space are simplified to only include essential features and meaningful actions. One might expect the agent to perform better in the symbolic version since it does not need to learn visuospatial and motor skills. However, the researchers found that current state-of-the-art RL methods do equally poorly in both versions, providing evidence that causal reasoning is a serious bottleneck in current RL methods. [13] is very similar to the symbolic version of *Alchemy* in that the agent can see the values of nodes of a causal graph, but not the edges in that graph. The agent can perform interventions by forcing a node to a particular value, and letting the node values be re-sampled. Both examples above involve two kinds of learning: an "inner loop" where the agent performs experiments to learn the causal structure that is specific to that episode, and an "outer loop" where the agent learns the ability to discover causal structures.

Finally, *ACRE* [52] is very similar to our approach in that it uses also uses the *blicket* detector paradigm and has essentially the same hypotheses space as ours. The main difference is that *ACRE* focuses on causal reasoning after the relevant observations have been generated by the experimenter, rather than letting an agent explore the causal relationships in an RL setting.

In conclusion, most environments are concerned with causal induction or generalization, rather than focusing on exploration [though see 48]. We instead focus on developing an environment with a controllable causal structure designed to allow us to measure the agents' ability to explore using causal overhypotheses. Moreover, none of these environments have been used with children. In fact, most of the environments discussed above either focus on causal knowledge that children are already familiar with in their daily lives (such as the laws of gravity), or are so complicated that children are not expected to perform well. Moreover, the complexity of those environments make it very hard to directly analyze and understand the exploration strategies of the agents, forcing us to only compare numeric metrics such as reward values. We aim to put agents and children in exactly the same environments, allowing us to use information about the real life causal exploration of these very effective child causal learners to inform agents.

Chapter 3

The Virtual “Blicket Detector” Environment

Although there is a large body of work on causal learning in children, and some experimental studies of active learning in the lab, as noted above, there are no studies analyzing children’s spontaneous actions as they freely explore the blicket machine, and none looking at how such exploration might lead to the learning of overhypotheses or how overhypotheses might influence learning.

We introduce a virtual, internet-hosted, representation of the standard blicket detector [16, 28] which is suitable not only for interaction with children, but also enables an *exact* comparison between children and reinforcement learning agents through an interface with



Figure 3.1: Screenshot of our virtual blicket detector. Children can interact with the blicket detector through a touch interface on an iPad, or through a point and click interface on any web-enabled machine.

OpenAI Gym [6]. In particular, this environment allows us to precisely record and analyze both children’s and agent’s actions. Unlike previous experiments [16, 28] which were mainly conducted in person, the internet-based environment also allows for a more diverse set of participants. A visualization of the online blicket detector is shown in 3.1.

Environment details Observations are a 3D rendering (256×256 RGB pixels) of the detector and a set of objects (3.1). In each episode, there are three objects, (which we refer to as A , B , and C for the rest of the report) and the blicket detector present in the environment. Each object in the environment has a unique color and shape. The environment is Markovian, as all objects and blicket detector states are visible at any given time. The children can put any combination of objects on the detector in any order, and click the check mark to test to see if its works or not. It lights up and makes a sound when the correct set of objects (i.e., blickets) are placed on top, and object combinations are permutation invariant in our setup (though children do not necessarily have this prior, see section 3.1). The required combinations of blickets to trigger the detector is governed by a particular causal structure that is randomly sampled per episode based on the condition. The action space consists of seven discrete actions. There are six actions for moving the objects (on/off for each of the three objects) and an additional action for pressing the check mark on the blicket detector, which checks if the existing objects make the blicket detector light up. Note, that the blicket detector will light up as long as a subset of the objects on the detector are blickets (e.g. putting three objects on the detector will also light up the detector).

Causal overhypotheses Our environment setup consists of a hierarchical causal structure, where the higher level structure is a causal hypotheses that determines the number of objects needed to light-up the blicket detector, and the lower level describes which particular objects are blickets. Similar to the setup in [28], we consider two causal hypotheses in our environment: *CONJUNCTIVE* and *DISJUNCTIVE*. In the *CONJUNCTIVE* condition, a pair of blickets must be on the machine (together) to activate it. In this case, A and B turn on the machine, but only when placed on the detector together so the only possible combinations that turn on the detector in the *CONJUNCTIVE* condition are AB , or ABC . The combinations that wouldn’t work are: A, B, C, AC and BC . In the other condition or *DISJUNCTIVE*, A and B are blickets which individually turn the blicket machine on. Therefore, the following combinations tun on the detector in the *DISJUNCTIVE* condition: A, AB, ABC, B, BC and AC .

3.1 Measuring Childrens’ Causal Exploration

We designed an experiment modeled on the blicket detector tasks [16] which allowed us to measure and analyze children’s exploration behavior and use of overhypotheses in an online environment that could also be given to an agent. We tested $N = 85$ children aged 4-5 years (20-23 children per condition) at a local science museum following IRB protocols. The exact

script we used can be found in the appendix of our paper <https://arxiv.org/abs/2202.10430>.

Conditions Our experimental setup consisted of a 2×2 design. Our first tier of conditions were CONJUNCTIVE and DISJUNCTIVE, reflecting different ways the detector could work, as described in chapter 3. Children also received one of two forms of evidence about the blicket detector, either suggesting (GIVEN HYPOTHESIS) or failing to suggest (NOT GIVEN HYPOTHESIS) a relevant hypothesis space.

Demonstration phase In both the DISJUNCTIVE and CONJUNCTIVE conditions, children first saw a video of a live demonstration of two different blicket machines. The machines had either a polka dot pattern or a stripe pattern and the objects varied in color and shape (sphere, pyramid and cube) as well, counterbalanced across conditions. The demonstrations for each machine and for each of the GIVEN HYPOTHESIS and NOT GIVEN HYPOTHESIS conditions are illustrated in 3.2. In the testing portion of the experiment, we specifically chose colors and shapes that were different from those used in the demonstration phase, so that children could not directly apply attribute-based overhypotheses from the demonstration phase to the exploration phase. In the GIVEN HYPOTHESIS condition the children received evidence that the blicket detectors could work in either a conjunctive or disjunctive way. In the NOT GIVEN HYPOTHESIS condition children only received ambiguous evidence about how the blicket machine might work, consistent with many overhypotheses. To avoid biasing the children, we did not give them any other evidence about how the machine works, beyond the demonstration video.

Exploration and test phase After they watched the demonstrations, the children were shown a new detector with a checkerboard pattern and ring, triangle and half dome objects (as in 3.1). They were told, “Look, I have a 3rd blicket detector. It could work like the polka dot one, or it could work like the striped one. Can you figure out how it works and which blocks make it go, and make the detector go yourself?” After making the machine light up the first time, the children were then asked “Great, is there something else you want to try?”. Once the child responded no to that question they were asked two final test questions. First, they were asked whether each object was or was not a blicket, and then they were asked how the machine worked.

Results

A major aim of this study was as a proof of concept that young children would actively and intelligently explore in this online environment and treat it as a causal system, as they do with the real-life blicket machines. We see this study as itself exploratory, but broadly, we hypothesized that children would exhibit systematic causal exploration rather than exploring randomly, and, as a result, make correct causal inferences.




Figure 3.2: Visualization of the demonstration the children see for which objects make the blicket detector light up, per condition, either Condition 1: GIVEN HYPOTHESIS or Condition 2: NOT GIVEN HYPOTHESIS.

Condition	# Participants	# Actions	# Combinations	Time (s)	Time to Success (s)
NOT GIVEN HYPOTHESIS (CONJUNCTIVE)	20	12.2 (8.19)	3.85 (2.59)	208.88 (78.64)	22.21 (19.73)
GIVEN HYPOTHESIS (CONJUNCTIVE)	22	12.68 (7.44)	5.36 (1.63)	164.89 (62.96)	23.61 (16.21)
NOT GIVEN HYPOTHESIS (DISJUNCTIVE)	20	8.95 (5.66)	3.7 (2.07)	181.87 (83.75)	12.35 (10.13)
GIVEN HYPOTHESIS (DISJUNCTIVE)	23	15.4 (8.39)	5.43 (1.34)	191.81 (63.81)	11.91 (20.76)

Table 3.1: Statistics of children’s exploration. Shown are the average number of checks (i.e., presses of the check mark) taken per condition, the average number of unique combinations tried per condition, the average time played per condition in seconds, and the average time played before seeing the blicket detector go on for the first time. Standard deviations are given in parenthesis.

Theory-driven exploration Following previous work on exploration in children [47, 46], we expected that children would not simply try to make the light go on, but that they would explore more extensively. As shown in 3.1, children in all conditions took less than 30 seconds to activate the detector for the first time but continued to explore for at least several minutes more. Moreover, children tried fewer unique combinations of objects than the total number of checks (i.e., the number of times they pressed the check mark to test the effect of the combination), indicating that children tested some combinations multiple times. Note that if a child clicked the check mark multiple times without any action in between, only the first time was counted and all subsequent checks were discarded. However, if the child performed any action, including taking all objects off and placing them back on in the same order, the second check would be included in our data. The number of combinations refers to how many different sets of objects the child tried, even if children placed the objects on the machine in different orders, which they frequently did. As we will discuss later in this chapter, the observation that children tried the same combinations multiple times and that they varied the order of the objects might indicate that children are considering additional

hypotheses which were not part of our initial analysis.

Inferential success Prior work has shown that children can make correct causal inferences about blicket machines given sufficiently informative data [28]. We similarly hypothesized that children would successfully be able to distinguish blickets from non-blickets given the evidence they generated during exploration. To test this, we compared how likely children were to report that true blickets are blickets (true positive rate) to how likely they are to report that non-blickets are blickets (false positive rate). The results are shown in 3.3 (left). Across all conditions, children were more likely to report that blickets were blickets than that non-blickets were blickets. While children were not perfectly able to determine the correct causal structure from the evidence they generated themselves, they showed some ability to do so across conditions. We did not ask the children to identify if the detector was conjunctive or disjunctive because it is challenging to get a meaningful answer on questions like this from children, especially in the NOT GIVEN HYPOTHESIS condition where they are not introduced to the concepts of conjunctive and disjunctive.

Effect of conjunctive vs. disjunctive [28] showed that, given sufficient evidence, children are equally good at making causal inferences about conjunctive and disjunctive structures. Similarly, we expected that children would be equally good at making inferences in the DISJUNCTIVE condition as the CONJUNCTIVE condition, and that they would exhibit similar amounts of exploration. To test this, we measured how long children explored (both in terms of time and number of actions taken) in the different conditions, as well as their success at discriminating between blickets and non-blickets.

As reported in Table 3.1, Children took approximately the same amount of time in both the CONJUNCTIVE (185.84 seconds) and DISJUNCTIVE (187.19 seconds) conditions. The same is true for the number of actions they performed (12.45 in CONJUNCTIVE and 12.39 in DISJUNCTIVE). Interestingly, even though it is easier to illuminate the blicket detector in the DISJUNCTIVE condition and children were faster at turning the blicket machine on for the first time in the DISJUNCTIVE condition, their exploration of both conditions was similar.

We also looked at how well children could discriminate between blickets and non-blickets across the different conditions. As illustrated in 3.3 (left), there is little difference between the conditions (with the exception of GIVEN HYPOTHESIS (DISJUNCTIVE), which we discuss further below). Taken together, these results indicate that the true causal structure of the blicket detector does not substantially influence how much children explore or how likely they are to come to a correct answer, consistent with earlier results [28].

Hypothesis space effects If an actor has evidence about the hypothesis space, the actor should be more efficient in the exploration of that space, and causal inference should be an easier task. Thus, we hypothesized that children in the GIVEN HYPOTHESIS conditions would explore less and be more accurate in their inferences than in the NOT GIVEN HYPOTHESIS condition.




Figure 3.3: Left: Children’s verbal replies about objects’ bliCKET-ness. Blue bars indicate the true positive rate, i.e. the proportion of times children said that an object was a bliCKET, given that it was a bliCKET. Yellow bars indicate the false positive rate, i.e. the proportion of times children said that an object was a bliCKET, given that it was not. The error bars indicate standard error. Random guesses would achieve 50% for both true positive rate and false positive rate. The differences between all the conditions are statistically significant with $p < 0.02$, except between the CONJUNCTIVE GIVEN HYPOTHESIS conditions. Right: The percentage of children who generated enough data to make a valid conclusion about which objects are bliCKETS (assuming an optimal inference procedure).




Figure 3.4: The average number of actions taken per condition, and the time that children require to make the first successful action (lighting up the detector) per causal model and condition.

First, to measure the amount of exploration that children performed, we looked at the amount of time they took and the number of actions they explored, as given in Table 3.1 and 3.4. When collapsing across causal structures, the results did not entirely align with our hypotheses: on average children tried more actions in a shorter time in the GIVEN HYPOTHESIS condition (178.65 seconds, 14.07 actions) than in the NOT GIVEN HYPOTHESIS condition (195.38 seconds, 10.58 actions).

Second, we examined if the data generated during exploration was sufficient to disambiguate the hypothesis space. 3.3 (right) indicates that children were substantially more likely to do so in the GIVEN HYPOTHESIS conditions (71% of the time) than in the NOT GIVEN HYPOTHESIS conditions (32% of the time, $p = 0.002$ using a standard 2-proportion Z test). Although children only tried a few more combinations more in the GIVEN HYPOTHESIS conditions, this exploration was more effective. This result is consistent with the hypothesis that without guidance, children may explore unbounded sets of overhypotheses, which may not be useful for discriminating between conjunctive and disjunctive structures.

Third, we looked at hypotheses that might naturally occur to children but should be ruled out in the GIVEN HYPOTHESIS condition. We identified 4 experiments that are indicative of these types of hypotheses. (1) No object needed. We counted the percentage of children who tested whether the machine would turn on without any objects on it. (2) Permutation. We counted the percentage of children who tried different orderings of the same set of objects (our environment preserves the order in which the the objects are put on top of the blicket detector). We only counted the cases where the same set of objects was tested successively. This makes it unlikely that, rather than trying to distinguish between permutations, the children simply forgot about the result of an experiment and wanted to test again. (3) Stochasticity. We counted the percentage of children who tested the same set of objects with the same permutation multiple times. Similar to (2), we only counted the cases where the same permutation was tested successively. (4) Anti-blicket. This refers to when two objects fail to turn the detector on, and the child chooses to take one object off and try again. In this hypotheses, the object being taken off is acting like an anti-blicket in that it prevents the machine from turning on. Note that this is only possible for the CONJUNCTIVE conditions. Since two out of three objects were always blickets, any two objects would always turn the machine on in DISJUNCTIVE conditions. Therefore, we only used the data from GIVEN HYPOTHESIS (CONJUNCTIVE) and NOT GIVEN HYPOTHESIS (CONJUNCTIVE) to calculate the percentages.

The percentages of children who performed these tests are reported in 3.2. All of the tests mentioned above are reasonable for children in the NOT GIVEN HYPOTHESIS condition; however, the children in the GIVEN HYPOTHESIS conditions could have ruled out all four. Thus it is surprising that children in both conditions tested all four hypotheses at similar rates. This suggests that children can be completely unaffected by prescribed knowledge when exploring hypotheses that are compelling to them. We hypothesize this to be a useful trait for causal discovery, and it remains an open question as to how AI systems can generate and judge these kinds of hypotheses.

Finally, we looked at whether children’s inferences about blickets and non-blickets were affected by being given the hypothesis space or not. As shown in 3.3 (left), there were generally no clear differences between conditions, again with the exception of the GIVEN HYPOTHESIS (DISJUNCTIVE) condition in which children were substantially more likely to correctly identify blickets. Although they were not more likely to generate sufficient data in this condition (3.3, right), they did try substantially more actions (Table 3.1).

Test	given hypothesis	not given hypothesis
No object needed	17.7%	17.6%
Stochasticity	51.1%	50%
Permutation	46.7%	52.5%
Anti-blicket	18.1%	20%

Table 3.2: Percentage of children who tested for hypotheses that should be ruled out in the GIVEN HYPOTHESIS conditions. If the children acted completely rationally, the first column should be 0 for all 4 hypotheses. Instead, we see the children are testing these hypotheses with near identical frequency regardless of the demonstration video they received

Summary Overall, our results suggest that children are able to effectively explore—particularly when given the relevant hypothesis space—and that they are often able to correctly identify blickets. However, children did not *always* generate sufficient evidence and could not perfectly discriminate between blickets and non-blickets (3.3). One explanation for this finding could be that the children were acting optimally (with noise); however, we favor another explanation: that children were optimizing for a wider range of alternative hypotheses. To differentiate between these explanations, in the next section we compare the results of children with several optimal models.

Chapter 4

Modeling Causal Learning Using Blickets

Given our experimental results, we aimed to characterize the optimal behavior for the blicket detector task, and used this optimal behavior to interpret the children’s decisions. As a first pass at understanding the motivations and priors of the children, we leveraged the relatively simple causal structure of the blicket environment to build several policies, which we used as a baseline for the children’s behavior. Our approach is similar to previous work [38, 36, 11] which has taken a Bayesian approach to optimal data selection and used such a model to describe adult behavior for various causal reasoning tasks. We used the unified environment described in Section 3, and the same experimental setup. Here, one action corresponds to putting any number of blickets on the detector and pressing the check button. For each action there are two possible observations (blicket machine turns on or remains off).

Policy Design In the blicket game, there are eleven possible causal structures, seven of which are disjunctive (*A-dis*, *B-dis*, *C-dis*, *AB-dis*, *AC-dis*, *BC-dis*, and *ABC-dis*, where “*A-dis*” refers to a disjunctive structure where *A* is a blicket and *B* and *C* are not), and four of which are conjunctive (*AB-con*, *AC-con*, *BC-con*, and *ABC-con*). We excluded degenerate causal structures such as *A-con* and *B-con* where no observations can distinguish between them, since conjunctive detectors require two blickets to light up. The goal of our models was to determine the causal structure through intervention in the environment. In this work, we investigated two policies through which we explored optimal behavior in the state space: a policy based on per-step information gain maximization (the *PER-STEP* model), and a policy which minimized the expected time to full disambiguation of the hypothesis space (the *MINIMUM-STEP* model).

per-step model The first policy optimizes for the expected per-action information gain, and aims to optimally discriminate between a set of fixed hypotheses through its actions by taking actions which minimize the uncertainty in the conditional posterior. We measured

uncertainty over all hypotheses using KL-divergence between the hypothesis distribution over causal hypotheses and some current prior distribution. More formally, let $h \in H$ be the possible hypotheses considered by this policy. The “usefulness” of a particular observation, o , resulting from an action a , is given by the difference between the posterior and prior distribution: $D_{\text{KL}}(p(h|o) \parallel p(h))$. The posterior probability for each individual hypothesis h is given by Bayes’ rule: $p(h|o) = p(o|h)p(h)/p(o)$. Before the first action, the prior distribution $p(o)$ is initialized using one of two possible choices (see below). For modeling multiple actions in sequence, the prior is defined to be to the posterior of the previous timestep, and actions are selected by maximizing the sum of the per-timestep divergences over the sequence. The PER-STEP model returns action sequences that are optimal for all possible ground-truth states of the blicket detector (i.e. A and B are blickets and the detector is conjunctive, A and C are blickets and the detector is disjunctive, etc.).

minimum-step model Instead of greedily optimizing for disambiguation as in the PER-STEP model, the second policy that we explored (the MINIMUM-STEP model) optimizes for the minimum number of steps required to fully disambiguate the hypothesis space. To determine this policy, we modeled the possible hypothesis space as a tree search problem, where each vertex represents the possible causal hypotheses, and each edge represents an action. The policy selects sequences of actions which minimize the expected depth of the tree rooted at each vertex. Unlike the PER-STEP model, this policy actively attempts to minimize the number of steps that are required for resolution.

Choice of prior We considered two possible prior beliefs for our models. The first prior, the UNIFORM prior, posits that that all causal structures are equally likely. However, in the GIVEN HYPOTHESIS condition, we implied equal likelihood between the CONJUNCTIVE and DISJUNCTIVE overhypotheses, and our experiments with children only contained causal structures with two blickets. Therefore, children could reasonably be expected to follow a different prior with three conjunctive structures and three disjunctive structures, are of which are equally likely. We encoded this possibility in a prior (the EXPERIMENTAL prior) which placed equal weight on conjunctive and disjunctive hypotheses.

Results, Analysis and Comparison to Children’s Exploration

The PER-STEP model and the MINIMUM-STEP models generated subtly different strategies for playing the game. In all experiments, the causal structure (or type of causal structure) was not known *a priori* and the policies had to determine both which causal structure (conjunctive vs. disjunctive) was present, and which objects in the scene were blickets. Because of the relatively small policy space, we could succinctly describe the policies in 4.1, 4.2, 4.3, and 4.4, to demonstrate the optimal strategies learned by each model. As expected, the MINIMUM-STEP agent optimizing for the minimum number of steps obtained a policy which is slightly better than PER-STEP, with MINIMUM-STEP achieving an expected time to full




Figure 4.1: The optimal policy for the PER-STEP model under UNIFORM prior. The observation tree is represented from left to right, with the first column representing the first action taken and the second column representing the resulting observation outcomes and subsequent action taken by the PER-STEP model. The model tests single objects, and if both tested objects are off, tries combinations of objects to attempt to reduce the uncertainty. The model is done when it knows the exact condition of the space.




Figure 4.2: The optimal policy for the PER-STEP model under the EXPERIMENTAL prior. This choice of prior reduces the number of policy branches, but also increases the average time to solution.




Figure 4.3: The optimal policy for the MINIMUM-STEP model under the UNIFORM prior.

Test Object C	(Detector on) - Test Object B	(Detector on) - Test Object A	(Detector off) Done (Detector on) Done (Detector on) Done (Detector off) Done (Detector on) Done (Detector off) Done (Detector on) Test Objects A,C		
Test Object C	(Detector off) - Test Object B	(Detector on) - Test Object A (Detector off) - Test Object A	(Detector off) Test Object A	(Detector on) Done (Detector off) Done (Detector on) Done (Detector off) Test objects A, C (Detector off) Test objects A, C	(Detector on) Done (Detector off) Done

Figure 4.4: The optimal policy for the MINIMUM-STEP model under EXPERIMENTAL prior.

information of 3.55 actions under the UNIFORM prior, and 3.50 actions under the EXPERIMENTAL prior, while PER-STEP achieved 3.72 expected actions to full information under the UNIFORM prior, and 4.0 expected actions under the EXPERIMENTAL prior. It is interesting to note that the PER-STEP model performed worse when given incorrect information (the EXPERIMENTAL prior), while the MINIMUM-STEP model is able to perform better in this scenario (the downside to the MINIMUM-STEP being the fact that it is usually intractable to compute in practice).

Next, we contrasted the models' behavior with that of the children. As shown in 4.5 (left), children always took more steps than the agents before they finished exploring the environment. Surprisingly, children did not show any notable difference between GIVEN HYPOTHESIS (CONJUNCTIVE) and GIVEN HYPOTHESIS (DISJUNCTIVE), whereas the agents differed significantly, with both optimal methods taking longer to explore for conjunctive spaces (note that the agents have no variance in expected time to goal, as these numbers are analytic). This increase in required actions for conjunctive causal structures demonstrates the bias towards disjunctive structures in the environment, as there are more possible disjunctive than conjunctive structures (7 vs. 4).

Even though children are explicitly encouraged to determine how the machine functions, while the agents always took enough actions in the space to fully disambiguate the causal structure, the children often did not take enough actions, especially when they were not given any information about the hypothesis space. 4.5 (right) shows the percentage of children who took enough actions to fully disambiguate the space. Unsurprisingly, when in GIVEN HYPOTHESIS, the children were far more likely to fully disambiguate the environment, demonstrating the power of fully guided exploration. There is little difference between the CONJUNCTIVE and DISJUNCTIVE conditions when it comes to exploration, suggesting that the children explored equally in both conditions, and did not favor disambiguation of conjunctive vs. disjunctive environments. This effect raises an interesting question: unlike RL agents, we do not have any explicit understanding of the reward function that children are optimizing during training time, and even though an effort was made to encourage the children to optimize the reward function of maximum information gain, the children may not optimized this reward alone. A key contribution of this work is to raise this question:




Figure 4.5: Left: The average number of combinations the children tried before they indicated they finished exploring, compared to the number of combinations the models needed to fully determine the ground truth. Agents use the EXPERIMENTAL prior. We find that in the CONJUNCTIVE GIVEN HYPOTHESIS condition the children’s average number of combinations is statistically significant compared to the MINIMUM-STEP model, but the difference between the PER-STEP model is not significant. In the DISJUNCTIVE GIVEN HYPOTHESIS condition the difference between the children and the MINIMUM-STEP and PER-STEP models is significant with $p < 0.0001$. Right: The percentage of children who generated enough observations to completely determine the ground truth in each condition. While conj. vs. disj. have no significant difference ($p > 0.9$), when given the hypothesis space, children are more likely to disambiguate the causal space ($p = 0.002$). The models always generate enough observations to disambiguate the causal space.

can we discover in future work what kinds of reward functions children optimize during overhypothesis discovery, and can we leverage these functions to build more powerful machine learning models?

Further analyses of children’s exploration As the previous results show, children’s exploration behavior was not well-captured by the proposed optimal models. By qualitatively examining children’s behavior in more detail, we observed that they appeared to be testing many different kinds of overhypotheses beyond the ground-truth conjunctive and disjunctive structures. 4.6 shows some example sequences of the combinations the children tried, indicating the various hypotheses they were testing. These qualitative examples suggest that children have a richer and larger hypothesis space than the models. For example, Participant 1 seemed to only consider the possibility that all three objects must trigger the detector, but attempted to determine if the order in which blickets are placed on the detector is important, as well as if the detector might be stochastic. Participant 2’s behavior was consistent with trying disjunctive conditions first, and then conjunctive conditions after that: a good example of systematic testing. Participant 3 only tested conjunctive behavior (with some testing of order). Participant 4 is another example where the child quickly discovered the




Figure 4.6: Example sequences of what objects children placed and tested on the blicket detector. The 0 column gives the ground truth blicket objects. Each row represents the order the blickets are placed on the detector, and the final row is green if the blicket detector is illuminated. The gray column represents when the child has finished exploring and is asked which objects are blickets.

causal structure, but then tested order and stochasticity.

In addition to causal structure, order, and stochasticity, childrens' overhypotheses may also take into account the attributes of objects, for example that blue objects are blickets or circular objects are blickets. Taking into account attributes would exponentially increase the size of the overhypothesis space, so we designed our experiment to try to minimize the likelihood that children would consider attribute-based overhypotheses. First, we used distinctly different colors and shapes in the exploration phase, compared to in the demonstration phase, so that children would not be able to directly apply attribute-based overhypotheses from the demonstration phase to the exploration phase. In addition, for each child we randomized which two of the three objects were blickets, in both the demonstration and exploration phases.

Unlike children, the optimal policies do not consider that the blicket detector may be stochastic or that there may be additional hypotheses outside of conjunctive and disjunctive. Determining a set of hypotheses that more accurately represents what the children actually consider would lead to a model which is more predictive of the children's behavior.

Chapter 5

Discussion & Conclusion

Exploration—and in particular, causal exploration—remains a challenge for modern RL agents. In this project, we investigated how human children engage in this type of exploration in order to draw insights that can be applied to artificial agents as well. To do so, we both introduced a novel “blicket” environment which enables testing exploration strategies of both humans and agents, and collected data on young children to observe how they explore in this domain. We found that children exhibited a diverse set of exploration strategies not well captured by naïve models optimizing for a small set of known overhypotheses. Notably, from our results it seems that children bring to bear extensive prior knowledge regarding how objects and mechanisms behave, and use this knowledge to generate a wide range of potential hypotheses.

These results have important implications for research aimed at developing RL agents. Most importantly, they suggest that effective exploration requires rich prior knowledge structured via causal overhypotheses. While in our experiments children’s behavior may technically be “suboptimal” in the context of simply disambiguating conjunctive vs. disjunctive hypotheses, in more realistic environments, children spend their time exploring an incredibly vast range of phenomena: block towers, bugs in the forest, digital devices, pets, social dynamics, and so on. To effectively explore in such a broad range of domains, children need to leverage rich prior knowledge; our experiments reflect that they do this by default even in simple settings. If we want agents that can perform a similar breadth of exploration, they similarly need to bring to bear rich causal knowledge about the world. Moreover, it may not be sufficient to have a large “bag-of-hypotheses”: like children, agents may need to organize their knowledge into hierarchical representations like overhypotheses in order to swiftly narrow in on the relevant subset of causal structures.

Although we have only presented results here for simple optimal models, these results pose a clear set of questions regarding RL, in particular: what is required for RL agents to exhibit exploration strategies similar to those children use? There are two challenges: what reward function should be used to train such agents, and under what training regime? A maximally “blank-slate” approach to our task might be: on each episode, sample whether the agent is in the CONJUNCTIVE or DISJUNCTIVE condition, and then reward it for making the

blicket detector turn on. Unfortunately, this training regime would not result in exploratory behavior at all: the optimal solution here would be to always just put all blocks on the detector. A more promising approach would be to reward agents for correctly identifying the causal structure, and optionally introduce downstream tasks that can only be accomplished with causal knowledge that is discovered during that episode. Some recent work has seen success using this approach [13, 34, 35]. But ultimately, this would only lead to behavior that converges to our optimal models in the ideal case. The children’s true motivation is likely a combination of the two factors mentioned above, and more carefully designed experiments would be needed to determine how they balances the two.

The second question requires more careful consideration. When presented with the blicket environment for the first time, children can use their prior knowledge and instincts to perform causal reasoning, essentially solving the task zero-shot. Since it is unclear how to teach an RL agent the abstract concept of causal learning without tethering it to a specific task or environment, it seems inevitable that the agent would need to be trained in a similar or identical environment using some distribution of causal hypotheses. The distribution of hypotheses an agent sees during training would have a strong influence on the exploratory behavior it can demonstrate at test time. One potential solution might be a “least knowledge” approach by enumerating the widest possible set of hypotheses and assigning them equal probability. However, this ignores the fact that children clearly have priors for what hypotheses are more likely. This intuitive judgement on priors is likely a key reason that makes children efficient at exploration, and it remains a challenge to measure those priors from noisy behavior.

We might also consider training agents on a broader set of overhypotheses (for example, that order matters, that the detector is stochastic, etc.), so that some portion of hypotheses, or even overhypotheses, can be held out during training. However, doing so can easily risk making the game too difficult or tedious for children. Since our goal is to learn from children’s exploration strategies, it is best to stay clear from the capacity of their cognitive abilities such as short term memory. But it is unclear what this set should be, or how to avoid it being overly task-specific, and we would need to be careful in how we categorize hypotheses into overhypotheses, as the criterion that seems most logical to the experimenter might not match the children’s intuition.

In conclusion, such “blank-slate” and “constructed” approaches do not appear consistent with how children learn about the world or how humans have developed over the course of evolution. We believe that developing environments and training regimes to tackle these questions is an essential aspect of future work in building agents that can efficiently and effectively explore.

Bibliography

- [1] Ossama Ahmed et al. “Causalworld: A robotic manipulation benchmark for causal structure and transfer learning”. In: *arXiv preprint arXiv:2010.04296* (2020).
- [2] Kelsey R. Allen, Kevin A. Smith, and Joshua B. Tenenbaum. “Rapid trial-and-error learning with simulation supports flexible tool use and physical reasoning”. In: *Proceedings of the National Academy of Sciences* 117.47 (2020), pp. 29302–29310. ISSN: 0027-8424. DOI: 10.1073/pnas.1912341117. eprint: <https://www.pnas.org/content/117/47/29302.full.pdf>. URL: <https://www.pnas.org/content/117/47/29302>.
- [3] Susan Amin et al. “A Survey of Exploration Methods in Reinforcement Learning”. In: *arXiv preprint arXiv:2109.00157* (2021).
- [4] Anton Bakhtin et al. “Phyre: A new benchmark for physical reasoning”. In: *Advances in Neural Information Processing Systems*. 2019, pp. 5082–5093.
- [5] Marc Bellemare et al. “Unifying count-based exploration and intrinsic motivation”. In: *Advances in neural information processing systems*. 2016, pp. 1471–1479.
- [6] Greg Brockman et al. *OpenAI Gym*. 2016. eprint: [arXiv:1606.01540](https://arxiv.org/abs/1606.01540).
- [7] Yuri Burda et al. “Exploration by random network distillation”. In: *arXiv preprint arXiv:1810.12894* (2018).
- [8] Maxime Chevalier-Boisvert et al. “BabyAI: A Platform to Study the Sample Efficiency of Grounded Language Learning”. In: *International Conference on Learning Representations*. 2018.
- [9] Karl Cobbe et al. “Quantifying generalization in reinforcement learning”. In: *arXiv preprint arXiv:1812.02341* (2018).
- [10] Karl Cobbe et al. “Quantifying generalization in reinforcement learning”. In: *International Conference on Machine Learning*. PMLR. 2019, pp. 1282–1289.
- [11] Anna Coenen, Jonathan D Nelson, and Todd M Gureckis. “Asking the right questions about the psychology of human inquiry: Nine open challenges”. In: *Psychonomic Bulletin & Review* 26.5 (2019), pp. 1548–1587.
- [12] Claire Cook, Noah D Goodman, and Laura E Schulz. “Where science starts: Spontaneous experiments in preschoolers’ exploratory play”. In: *Cognition* 120.3 (2011), pp. 341–349.

- [13] Ishita Dasgupta et al. “Causal reasoning from meta-reinforcement learning”. In: *arXiv preprint arXiv:1901.08162* (2019).
- [14] Alison Gopnik. “Scientific thinking in young children: Theoretical advances, empirical research, and policy implications”. In: *Science* 337.6102 (2012), pp. 1623–1627.
- [15] Alison Gopnik and David M Sobel. “Detectingblickets: How young children use information about novel causal powers in categorization and induction”. In: *Child development* 71.5 (2000), pp. 1205–1222.
- [16] Alison Gopnik and David M. Sobel. “Detectingblickets: how young children use information about novel causal powers in categorization and induction”. In: *Child Development* 71.5 (2000), pp. 1205–1222.
- [17] Alison Gopnik and Henry M Wellman. “Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory.” In: *Psychological bulletin* 138.6 (2012), p. 1085.
- [18] Alison Gopnik et al. “A theory of causal learning in children: causal maps and Bayes nets.” In: *Psychological review* 111.1 (2004), p. 3.
- [19] Alison Gopnik et al. “Causal learning mechanisms in very young children: two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation.” In: *Developmental psychology* 37.5 (2001), p. 620.
- [20] Alison Gopnik et al. “Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood”. In: *Proceedings of the National Academy of Sciences* 114.30 (2017), pp. 7892–7899.
- [21] Thomas L Griffiths and Joshua B Tenenbaum. “Theory-based causal induction.” In: *Psychological review* 116.4 (2009), p. 661.
- [22] Pim de Haan, Dinesh Jayaraman, and Sergey Levine. “Causal confusion in imitation learning”. In: *Advances in Neural Information Processing Systems* 32 (2019), pp. 11698–11709.
- [23] Stephen James et al. “Rlbench: The robot learning benchmark & learning environment”. In: *IEEE Robotics and Automation Letters* 5.2 (2020), pp. 3019–3026.
- [24] Nan Rosemary Ke et al. “Systematic Evaluation of Causal Discovery in Visual Model Based Reinforcement Learning”. In: *arXiv preprint arXiv:2107.00848* (2021).
- [25] Charles Kemp, Andrew Perfors, and Joshua B. Tenenbaum. “Learning overhypotheses with hierarchical Bayesian models”. In: *Developmental Science* 10.3 (2007), pp. 307–321.
- [26] Tamar Kushnir and Alison Gopnik. “Conditional probability versus spatial contiguity in causal learning: Preschoolers use new contingency evidence to overcome prior spatial assumptions.” In: *Developmental psychology* 43.1 (2007), p. 186.

- [27] Christine H. Legare. “Exploring explanation: explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children”. In: *Child Development* 83.1 (2012), pp. 173–185.
- [28] Christopher G Lucas et al. “When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships”. In: *Cognition* 131.2 (2014), pp. 284–299.
- [29] Marlos C Machado, Marc G Bellemare, and Michael Bowling. “Count-based exploration with the successor representation”. In: *arXiv preprint arXiv:1807.11622* (2018).
- [30] Marlos C Machado et al. “Revisiting the arcade learning environment: Evaluation protocols and open problems for general agents”. In: *Journal of Artificial Intelligence Research* 61 (2018), pp. 523–562.
- [31] Jarryd Martin et al. “Count-based exploration in feature space for reinforcement learning”. In: *arXiv preprint arXiv:1706.08090* (2017).
- [32] Daniel McDuff et al. “CausalCity: Complex Simulations with Agency for Causal Discovery and Reasoning”. In: *arXiv preprint arXiv:2106.13364* (2021).
- [33] Andrew N Meltzoff, Anna Waismeyer, and Alison Gopnik. “Learning about causes from people: observational causal learning in 24-month-old infants.” In: *Developmental psychology* 48.5 (2012), p. 1215.
- [34] Vladimir Mikulik et al. “Meta-trained agents implement Bayes-optimal agents”. In: *arXiv preprint arXiv:2010.11223* (2020).
- [35] Suraj Nair et al. “Causal induction from visual observations for goal directed tasks”. In: *arXiv preprint arXiv:1910.01751* (2019).
- [36] Jonathan D Nelson et al. “Experience matters: Information acquisition optimizes probability gain”. In: *Psychological science* 21.7 (2010), pp. 960–969.
- [37] Alex Nichol et al. “Gotta learn fast: A new benchmark for generalization in rl”. In: *arXiv preprint arXiv:1804.03720* (2018).
- [38] Mike Oaksford and Nick Chater. “A rational analysis of the selection task as optimal data selection.” In: *Psychological Review* 101.4 (1994), p. 608.
- [39] Ian Osband et al. “Deep exploration via bootstrapped DQN”. In: *Advances in neural information processing systems* 29 (2016), pp. 4026–4034.
- [40] Georg Ostrovski et al. “Count-based exploration with neural density models”. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR.org. 2017, pp. 2721–2730.
- [41] Charles Packer et al. “Assessing generalization in deep reinforcement learning”. In: *arXiv preprint arXiv:1810.12282* (2018).
- [42] Deepak Pathak, Dhiraj Gandhi, and Abhinav Gupta. “Self-supervised exploration via disagreement”. In: *arXiv preprint arXiv:1906.04161* (2019).

- [43] Deepak Pathak et al. “Curiosity-driven exploration by self-supervised prediction”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 2017, pp. 16–17.
- [44] Danilo J Rezende et al. “Causally correct partial models for reinforcement learning”. In: *arXiv preprint arXiv:2002.02836* (2020).
- [45] Jürgen Schmidhuber. “Curious model-building control systems”. In: *Proc. international joint conference on neural networks*. 1991, pp. 1458–1463.
- [46] Laura E. Schulz. “The origins of inquiry: inductive inference and exploration in early childhood”. In: *Trends in Cognitive Sciences* 16.7 (2012), pp. 382–389.
- [47] Laura E. Schulz and Elizabeth Baraff Bonawitz. “Serious fun: preschoolers engage in more exploratory play when evidence is confounded”. In: *Developmental Psychology* 43.4 (2007), pp. 1045–1050.
- [48] Sumedh A Sontakke et al. “Causal curiosity: RL agents discovering self-supervised experiments for causal representation learning”. In: *International Conference on Machine Learning*. PMLR. 2021, pp. 9848–9858.
- [49] Haoran Tang et al. “# exploration: A study of count-based exploration for deep reinforcement learning”. In: *Advances in neural information processing systems*. 2017, pp. 2753–2762.
- [50] Jane X. Wang et al. “Alchemy: A structured task distribution for meta-reinforcement learning”. In: *arXiv preprint arXiv:2102.02926* (2021).
- [51] Tianhe Yu et al. “Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning”. In: *arXiv preprint arXiv:1910.10897* (2019).
- [52] Chi Zhang et al. “Acre: Abstract causal reasoning beyond covariation”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, pp. 10643–10653.