# A Computational Characterization of Domain-Based Causal Reasoning Development in Children

by

Caroline Bradley Aronoff

B.S., Massachusetts Institute of Technology (2015)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2018



Author . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
Department of Electrical Engineering and Computer Science
May 24, 2018

Certified by . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
Patrick H. Winston
Ford Professor of Artificial Intelligence and Computer Science
Thesis Supervisor

Accepted by . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
Katrina LaCurts
Chairman, Masters of Engineering Thesis Committee

# A Computational Characterization of Domain-Based Causal Reasoning Development in Children

by

Caroline Bradley Aronoff

Submitted to the Department of Electrical Engineering and Computer Science
on May 24, 2018, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

## Abstract

To better understand human intelligence, we must first understand how humans use and learn from stories. One important aspect of how humans learn from stories is our ability to reason about cause and effect.

Psychological evidence suggests that when children develop the ability to learn cause-and-effect relationships from stories, they do so in discrete stages where each new stage enables the child to incorporate new kinds of information. In this thesis, I attempt to shed light on the mechanisms that underlie the development of causal reasoning in children. I create a behavior-level model, an explanatory theory, and an explanation-level model that account for the developmental stages. I implement these models on top of the Genesis Story Understanding System. The result is a *psychologically plausible* explanation-level model that captures the observed causal reasoning behaviors of children at different stages of developments. The model also takes the observations from psychological evidence to another level by proposing mechanisms that enable such development in children.

Thesis Supervisor: Patrick H. Winston
Title: Ford Professor of Artificial Intelligence and Computer Science

# Acknowledgments

I would be remiss to not take the time and space here to thank all of the people without whom this endeavor would not have been possible:

First and foremost, I would like to thank **Patrick Winston** for giving me the support and guidance I needed to make this thesis a reality. Thank you for sharing your knowledge, wisdom, and passion for Artificial Intelligence. And thank you for helping me take my wildest ideas and somehow turn them into an actual project. It has been an honour and a pleasure to be a part of your research group.

Thank you **Dylan Holmes** for the countless hours of helping me brainstorm, debug, and edit all aspects of this thesis. You always made time to help me when I asked, whether it was for my research or for my other story telling endeavors.

Thank you **Laura Schulz** for providing invaluable insight into your work and the many recommendations of other relevant studies to read. Without your guidance, my work would not have been nearly as well based in cognitive science.

Thank you to all my friends and colleagues in the **Genesis Group**—especially Suri, Ada, Dylan, Jessica, and Jake—for providing support, encouragement, and consistent fascinating discussions about both human and artificial intelligence.

Thank you to all the women in **Alpha Epsilon Phi - Beta Epsilon Chapter** who have shown me nothing but love and support throughout this entire process.

Thank you **Nanu Roitman**. You have been my rock throughout this crazy journey. Thank you for believing in me and for always being a voice of reason and sanity when I needed one.

Finally, thank you to **my parents** for always encouraging me to pursue my interests and passions in life and for never wavering in their belief that I would succeed in anything I put my mind to.

# Contents


**1 Introduction** **11**

1.1 Vision . . . 11

1.1.1 Why Causal Reasoning? . . . 12

1.1.2 Psychological Plausibility . . . 13

1.1.3 A Model for Computational Psychology . . . 15

1.2 Approach . . . 15

1.2.1 Why stories? . . . 16

1.2.2 Why Domain-Specific Causal Reasoning? . . . 16

1.3 Novel abilities of Genesis (as of May 2018) . . . 17

1.4 Setting the Stage . . . 17

**2 Can Being Scared Cause Tummy Aches?** **19**

2.1 Extension to this Thesis . . . 23

**3 Genesis System Background** **25**

3.1 The Rules . . . 27

3.1.1 Prediction Rules . . . 27

3.1.2 Explanation Rules . . . 28

3.2 Concept Patterns . . . 29

**4 Building The Computational Models** **31**

4.1 The Behavior-Level Model . . . 32

4.1.1 Three-year-old . . . 32

4.1.2 Four-to-five-year-old . . . . . . . . . . . . . . . . . . . . . . 34
4.1.3 Three-and-a-half-year-old . . . . . . . . . . . . . . . . . . . . 36
4.2 The Psychologically Plausible (Explanation-Level) Model . . . . . . . 38
4.2.1 Discussion of Psychologically Plausible Models Considered . . 38
4.2.2 Representational Choice . . . . . . . . . . . . . . . . . . . . 46
4.2.3 Rules for the Model . . . . . . . . . . . . . . . . . . . . . . 46
4.2.4 Concept Patterns . . . . . . . . . . . . . . . . . . . . . . . 50
4.3 Converting the experiment to be Genesis Compatible . . . . . . . . . 54
4.3.1 The stories . . . . . . . . . . . . . . . . . . . . . . . . . . 54
4.3.2 The questions . . . . . . . . . . . . . . . . . . . . . . . . . 56
4.4 High-Level Structure . . . . . . . . . . . . . . . . . . . . . . . . . 56
4.4.1 Expert Class and Wiring . . . . . . . . . . . . . . . . . . . . 57
4.4.2 Filtering and Answering Questions . . . . . . . . . . . . . . . 57

**5 Contributions** **59**

**A Stories** **61**

# List of Figures

2-1 Experimental results from Schulz experiment 1 . . . . . . . . . . . . . . 22
2-2 Analysis of Mechanisms and Behavior . . . . . . . . . . . . . . . . . . . 24
3-1 Macbeth Elaboration Graph . . . . . . . . . . . . . . . . . . . . . . . . 26
3-2 Prediction Rule Example Text . . . . . . . . . . . . . . . . . . . . . . . 28
3-3 Prediction Rule Elaboration Graph Example . . . . . . . . . . . . . . . 28
3-4 Explanation Rule Example Text . . . . . . . . . . . . . . . . . . . . . . 29
3-5 Explanation Rule Elaboration Graph Example . . . . . . . . . . . . . . 29
3-6 Concept Pattern Example Text . . . . . . . . . . . . . . . . . . . . . . 30
3-7 Concept Pattern of Revenge . . . . . . . . . . . . . . . . . . . . . . . . 30
4-1 Three-Year-Old Within-Domain Results . . . . . . . . . . . . . . . . . 33
4-2 Three-Year-Old Cross-Domain Results . . . . . . . . . . . . . . . . . . 34
4-3 Four-to-Five-Year-Old Within-Domain Results . . . . . . . . . . . . . . 36
4-4 Four-to-Five-Year-Old Cross-Domain Results . . . . . . . . . . . . . . . 36
4-5 Three-and-a-half-Year-Old Within-Domain Results . . . . . . . . . . . 37
4-6 Three-and-a-half-Year-Old Cross-Domain Results . . . . . . . . . . . . 37
4-7 Three-year-old Abstract Representation of Domain Knowledge . . . . 42
4-8 Three-and-a-half-year-old Abstract Representation of Domain Knowledge . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 43
4-9 Four-to-five-year-old Abstract Representation of Domain Knowledge . 44
4-10 Examples of Prediction Rules Used in Explanation Level Models . . . 47
4-11 Examples of Explanation Rules in Explanation Level Models . . . . . 48

4-12 Examples of Explanation Rules in Four-to-Five-Year-Old Explanation Level Model . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 49
4-13 Three and Three-and-a-half-Year-Old Rule Based Elaboration Graph 49
4-14 Four-to-Five-Year-Old Rule Based Elaboration Graph . . . . . . . . . 50
4-15 Concept Patterns for Sickness . . . . . . . . . . . . . . . . . . . . . . 51
4-16 Three-year-old Concept Pattern Elaboration Graph . . . . . . . . . . 52
4-17 Four-to-five-year-old Concept Pattern Elaboration Graph Psychological Cause . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 53

# Chapter 1

# Introduction

## 1.1 Vision

If we are to understand human intelligence, we must first understand how we humans understand and reason about stories. Stories and symbolic knowledge are a distinctive and essential part of human intelligence. Dating back nearly 35,000 years ago, there is evidence of symbolic representations in the form of cave paintings [8]. These paintings are the earliest that have been found, and the fact that they exist shows that the building blocks of symbolic thought were in place. These building blocks enabled the process of creating lasting ways of conveying information; a form of telling stories through symbols. This uniquely human trait of symbolic knowledge and representation is, in essence, a means of telling stories to ourselves and to each other. In order to truly understand human intelligence, we must first understand how humans use and learn from stories [15]. Because an essential part of story understanding is causal reasoning, we must understand how we reason about cause and effect.

In this thesis, I propose a cognitive theory and corresponding humanly-plausible model of how a particular aspect of causal reasoning develops in children.

### 1.1.1 Why Causal Reasoning?

Causal reasoning and understanding is a necessary skill in daily life. There is often sparse information presented for which we need to provide the missing links. For example, answering why and how questions like, "why is my friend upset?" requires causal reasoning. When examining a patient, doctors are often searching for an explanation or cause of the symptoms that the patient presents. Detectives are required to piece together the causes and chains of events within crime scenarios. When we read, we are often called upon by authors to infer how one event leads to another in order to understand stories. And children are constantly confronted with new world experiences through which they need to learn causal relations in order to understand how our world works. Causal reasoning is an absolutely essential skill that we use to understand the world around us.

We can better understand causal reasoning by understanding how it develops in children. This way, rather than study the fully developed behavior, we can study the building blocks, the underlying mechanisms, as they develop. There is significant interest among psychologists and cognitive scientists in trying to understand how causal reasoning develops. One way that Gweon and Schulz studied the development of causal reasoning is by investigating whether children can infer causes based upon evidence they have seen [4]. They found that children as young as 16 months old can use some basic statistical evidence in order to infer the cause of an event.

However, as the causal reasoning task becomes more complex this ability disappears [10]. For example, when domain-specific knowledge is introduced into causal inference scenarios, three-year-old children ignore the statistical evidence presented to them [10]. Given this information, it stands to reason that there must be something more than just statistical calculations alone going on in our brains when we are tasked with causal reasoning deductions. Moreover, while one may be still be able to model this behavior with a statistical model, I think it is important to question whether children, or humans in general, are best modeled as a form of statistician. With this in mind, I believe that it is important that, when modeling causal reasoning behaviors, the underlying model is more than a just statistical calculation in order to make the model meaningful.

It is with this philosophy in mind that my thesis was developed:

In order to understand human intelligence, we must first understand how we develop our causal reasoning capabilities; a good first step towards this understanding is developing a psychologically plausible model for this behavior. The model must not only describe behavior, but explain it.

### 1.1.2 Psychological Plausibility

In order for a computational model to be "psychologically plausible", the model must be:

1. Explanatory
2. Compositional
3. Sensible

These three traits, as I describe below, capture important criteria for a theory to be a "good explanation" [11] of a phenomenon.

#### Explanatory

When I say that a model should be explanatory, I mean that the model should be able to, in some manner, describe how it came up with the results that it produced. For example, if a doctor were to diagnose a patient, they would be able to describe the steps that they took in order to produce the diagnosis. An inferior model would simply produce the same diagnosis as the doctor in this example, but a better model would be able to explain itself the way people do. Even when people don't really know the answer to a causal reasoning question, they will take guesses and explain themselves [7]. In order to reach human-like intelligence, our models should have the same capability.

### Compositional

Compositionality helps support the explainability of a system. Essentially, a compositional model would be made up of different, well defined parts, determine how they fit together, and how the behavior of the whole emerges from the behavior of the individual pieces. If we have a model that exhibits these properties of compositionality, then we can test the individual parts of the model as well as the whole.What I mean to say is that a compositional model supports a divide and conquer approach to testing a theory. This kind of a model is necessary in order to be psychologically plausible because these models are trying to describe a theory of a mechanism for what might be happening in our minds, rather than just a description of the resulting behavior. Descriptive models may produce correct behavior, but they do not shed any light on what is happening under the hood.

### Sensible

Essentially, what I mean by a sensible model here is that the model must not only produce the desired behavior, but it should also suggest a scientifically plausible mechanism for the behavior or how that behavior emerges from a particular architecture or well-understood parts. By a scientifically plausible mechanism, I mean to say that the model operates under humanly-realistic constraint. Model building should not simply be an exercise in engineering, but should be a scientific endeavor to create a model that explains a phenomenon.

This goes hand in hand with the explanatory requirement. A psychologically plausible model should be a model designed with some backing in a cognitive, psychological, or otherwise relevant scientific theory. By having a design driven by solid scientific theories, we can produce a model that is scientifically relevant, explainable, and falsifiable. A good model should not only produce proper answers where humans succeed, but should also fail where humans fail. In this sense it somewhat relates to being explanatory, so that we can see if it fails in the same way(s) as humans do for a particular task. It is in the realm of sensibility that purely statistical models run into some problems. While statistical models may be able to capture the desired behavior, they do little to explain what is happening [3], and they often exhibit super-human capabilities.

In summation, these are the three traits required for a model to be psychologically plausible.

In this thesis, I have developed a model that is psychologically plausible in that it is exaplanatory, compositional, and sensible.

### 1.1.3 A Model for Computational Psychology

I intend for this thesis to serve as one possible model for how to create computational models of psychological phenomena.

The main point I would like to make here is that any good computational model of a psychological phenomenon (or even for a more broad spectrum of scientific phenomena) should be more than just a system that produces the same or similar behavior. It should be a system that also strives to function in an analogous way to the phenomema you are trying to model, or the theory you are trying to model. In essence, any good computational model for these endeavors should try to encompass the three properties I discussed as part of a psychologically plausible model. I believe that this a good way to structure and design scientifically sound computational models.

## 1.2 Approach

In this thesis, I tackle the problem of how children develop causal reasoning capabilities and why they initially fail when that reasoning involves domain-specific knowledge. Using a story-based framework, I recreate aspects of a psychological study by Schulz et al. [10] and take it to another level with a computational implementation that is explanatory, compositional, and sensible.

### 1.2.1 Why stories?

In order to understand human intelligence we must understand how humans tell, understand, and recombine stories because the mechanisms that enable these abilities are what separate human intelligence from that of other primates [15]. Humans, as opposed to other primates, have the ability to perform what Noam Chomsky calls a *merge* operation which is enabled by a completed anatomical loop in our brains that other primates do not have [14]. This *merge* mechanism is what allows us to create complex, highly nested symbolic representations of various sorts, which directly gives rise to our story related abilities [14]. Moreover, stories are an integral part of how we learn. We learn about morals from fairy tales, history through anecdotes, and even things like recipes can be thought of as a special sort of story. Furthermore, causal reasoning is an important aspect of story understanding. Therefore, it stands to reason that in order to understand causal reasoning development it is reasonable to investigate it through the lens of story understanding.

Additionally, as a representation, stories have important properties that other representations lack. Stories are explanatory, compositional, and easy for people to understand. They are easily crafted to suit the purposes and needs of any given study which makes them particularly useful for studying the relationship between domain knowledge, causal reasoning capabilities, and statistical data integration. Finally, practically, the Genesis Story Understanding System, on which this thesis is based, provides a powerful framework for story based operations such as pattern matching and incorporating commonsense knowledge.

### 1.2.2 Why Domain-Specific Causal Reasoning?

Out of all the studies done on causal reasoning, why choose to emulate this study ([10]) in particular?

The Schulz et al. paper [10] distinguished very well between implementation and theory. This separation allowed me to keep the theoretical aspect sound, while modifying the implementation as necessitated by the constraints of the system I was using (the Genesis system). Essentially, while I tried to keep the surface level methods of my thesis as close to those in the study as possible, the flexibility in the design of the methods of the original study allowed me to take liberties as needed without compromising the overall idea or design of the experiment. The fact that the study used stories as its main experimental component was what drew me to the study, and the fact that the stories are easily modifiable without losing their meaning or intended design is what cemented its place in the center of my thesis.

Additionally, practically speaking, this was a study that was computationally tractable. The Genesis system for story understanding already existed, and the study was entirely based upon children understanding stories. Essentially, after reading the study, I knew that it would be a reasonable and interesting endeavor to attempt to create computational models that leverage the Genesis system in order to reproduce the results of this psychological study.

## 1.3 Novel abilities of Genesis (as of May 2018)

Through my work, the Genesis system can now:

- Answer causal reasoning questions
- Deploy probabilistic commonsense inferences
- Model causal reasoning capabilities of different age groups
- Handle domain-related knowledge

## 1.4 Setting the Stage

In the rest of this thesis, I describe my computational, story-based model which captures the development of domain-based causal reasoning in children. In Chapter 2, I give an overview of the psychological study that I based my thesis upon. In Chapter 3, I provide an overview of the Genesis Story Understanding System with a focus on some features that were particularly relevant to my work. Having provided the necessary background information, motivation, and theories, in Chapter 4, I discuss the models themselves. I start with a behavior level model which replicates the behavior in the study. I then define and evaluate five alternative models which are more psychologically plausible. Finally, I describe the implementation of my explanation level model. In Chapter 5, I summarize my overall contributions.

# Chapter 2

# Can Being Scared Cause Tummy Aches?

To fully understand what I contribute in this thesis, you must first know a bit about the study I used as the basis for my theory and models: "Can Being Scared Cause Tummy Aches? Naive Theories, Ambiguous Evidence, and Preschoolers' Causal Inferences" [10]. There were three different experiments detailed in the study, and my thesis is focusing explicitly on the first of these 3 experiments. The purpose of the study was to investigate how children's general-purpose statistical reasoning and domain-specific causal learning abilities interact. Essentially, there are domain-general learning strategies where children learn from statistical evidence and domain-specific theories that constrain children's beliefs about what may be possible outcomes of a domain-related event.

To begin, I shall define some terms: within-domain and cross-domain. For the purposes of this thesis, I define a domain to be a conceptual category of information. For example, physical actions (running, playing, contacting, etc.) might constitute one domain and psychological actions (thinking, feeling, etc.) might constitute another. These domains contain distinct sets of concepts. When I say "within-domain" what I mean is that all the relevant actions come from the same conceptual category (i.e., all physical *or* all psychological) as opposed to "cross-domain" where the relevant actions come from a mixture of conceptual categories (i.e., a mixture of physical *and* psychological).

In order to test how general-purpose statistical reasoning and domain-specific causal learning abilities interact, the experimenters read each child two stories, one within-domain story and one cross-domain story, and then asked questions of the form "Why does X? Is it because of Y or Z?". Each story had one event that consistently co-occurred with the effect X and multiple confounding events that also inconsistently co-occurred with X; these events are represented by my placeholders Y and Z in that question. For the within-domain story, all of X, Y, and Z were in the same (physical) domain, while in the cross-domain story one of Y and Z was a domain-inappropriate (psychological) potential cause while the other was in the same (physical) domain as X. As mentioned earlier, in each story one of the potential causes Y and Z always preceded effect X. In the cross-domain story, the potential cause that always preceded the effect was the domain-inappropriate cause (or cross-domain cause).

Additionally, half the children in the study were part of the control or "baseline" group. These children were only read two scenes from each story, as opposed to the children in the experimental or "evidence" group who were read the entirety of each story. This allowed the experimenters to understand and note any intrinsic biases that the children may have in each scenario.

These stories would easily provide enough statistical evidence for an adult to be able to confidently assert which of the potential causes in question was the true cause of the effect in question, but children's abilities to reason statistically with domain-specific knowledge seems to be limited. This ability seems to develop over the course of approximately a year of age. This study looked at three different age groups: three-year-olds, three-and-a-half-year-olds, and four-to-five-year-olds, each of which had distinct behavior.

What this experiment showed is that these age groups handle the integration of statistical evidence with domain-specific beliefs in systematically different ways. The three-year-olds apparently failed to learn from the evidence in either scenario while the 4-to-5-year-olds were able to learn from evidence in both scenarios, albeit with some more difficulty over all in the cross-domain task. What is particularly interesting though, is that the three-and-a-half-year-olds were able to learn from the evidence in the within-domain scenario, but apparently unable to do so in the cross-domain scenario.

These results are shown in Figure 2-1 [10]. The black portion of the bars in this Figure represents the number of children in that age group who chose the "correct" answer, referred to in the Schulz et al. paper as 'A'. These results show that those children in the baseline condition were all decently biased against 'A' (the cross-domain cause) in the cross-domain task, while they seemed to answer at chance for the within-domain scenario.

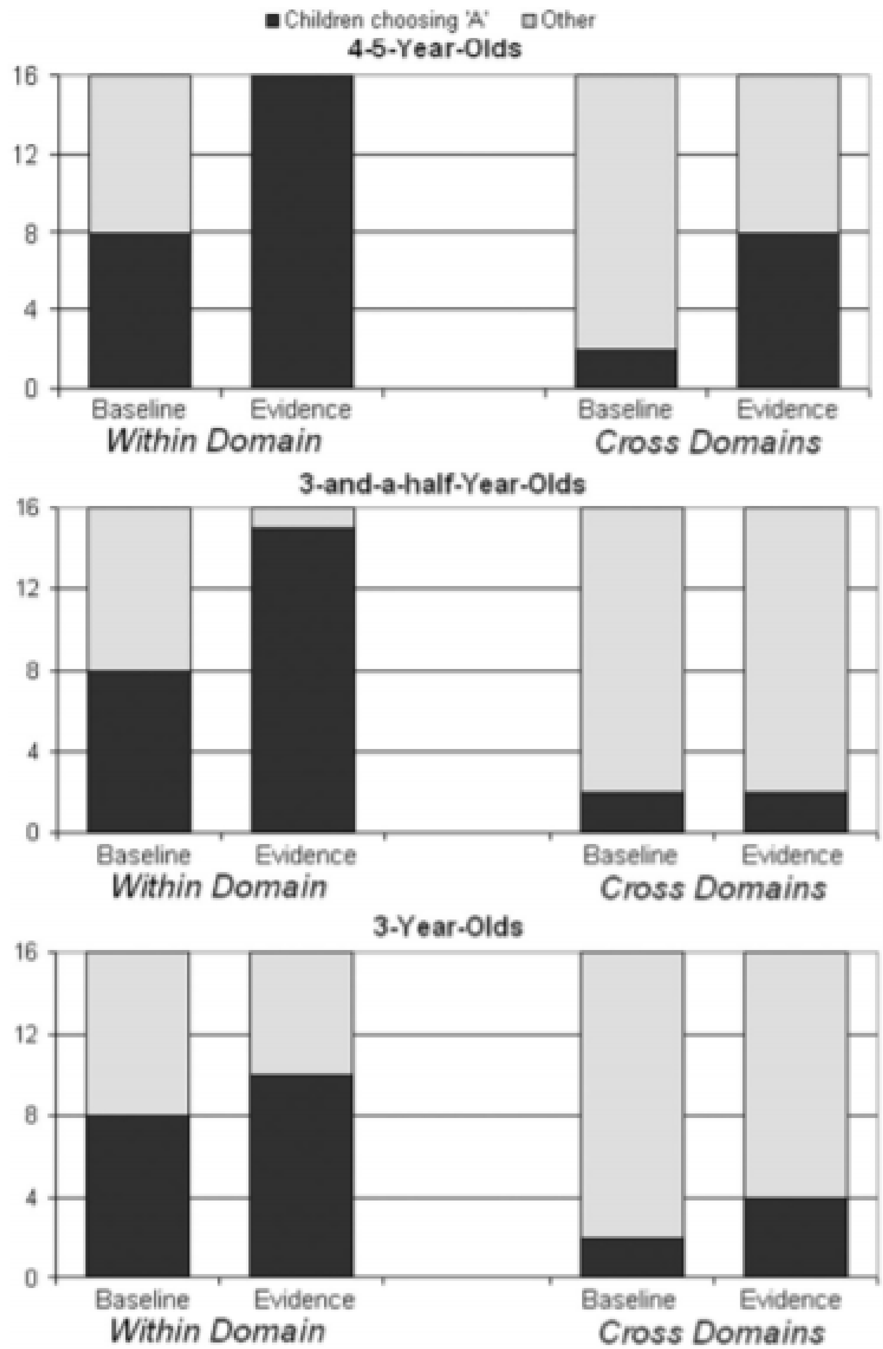


*Figure 2.* Children's responses to the storybook task in Experiment 1. The vertical axis shows the number of children selecting the different responses.

Figure 2-1: Children of different ages exhibit differing ability to integrate domain-related and statistical information according to the Schulz et al. study [10].

## 2.1 Extension to this Thesis

The main contribution of my thesis is the development of an explanatory model for the observations that Schulz et al. make. Rather than providing a statistical description of the behavior, my model leverages story understanding capabilities such that it is explanatory, compositional and sensible. The authors of the Schulz et al. study offered three possible explanations as to why this behavior occurred. One of the most striking explanations is that younger children "... might have difficulty making inferences from ambiguous statistical data...the ability of the 3.5-year-olds to interpret data of this complexity is fragile, any increase in task difficulty... might compromise children's ability to evaluate the evidence" [10]. For this thesis, I have taken this hypothesis and expanded upon it.

I propose an abstract theory about how children might develop the ability to reason about cross-domain events. The authors suggest that an increase in task difficulty from within-domain to cross-domain reasoning tasks may be the cause of the three-and-a-half-year-olds ability to use evidence in the within-domain case, but not in the cross-domain case. I am proposing a theory about what might be happening between the three-and-a-half-year-olds and four-to-five-year-olds to lessen the difficulty of the cross-domain task.

I hypothesize that the three and three-and-a-half-year-olds, abstractly speaking, may handle information about different domains separately from one another. You can visualize this as having a Venn diagram consisting of separate circles each containing the information about a different domain (Figure 4-7 seen in Chapter 4). Over time, the knowledge within each of these circles (each of the domains) becomes more interconnected, which would make it easier for children to reason about within-domain causal relations (Figure 4-8 seen in Chapter 4). This could account for the gap between three and three-and-a-half-year-old behavior. Furthermore, as we continue to grow and learn, the knowledge in circles may start to form connections with information from a different circle (cross-domain connections) that would make it easier to reason about cross-domain causal relations (Figure 4-9 seen in Chapter 4).

This could account for the gap in reasoning behavior between three-and-a-half and four-to-five-year-olds. This theory can be thought of, speculatively, as a neuronally inspired mechanistic approach that proposes that the development of a cognitive ability is due to the creation of new "neuronal" links between information-specific regions [2].

Even more abstractly, we can think of these two abilities (statistical information reasoning and cross-domain understanding) as two switches that "turn on" over time (Figure 2-2). This conceptual paradigm is better shown through Figure 2-2.

| | | Cross-Domain Understanding | |
| --- | --- | --- | --- |
| | | False | True |
| Integration of Statistical Information | False | 3-year-old | *3-year-old savant |
| | True | 3.5-year-old | 4-year-old |

Figure 2-2: Analysis of Mechanisms and Behavior: This table shows how the combinations of behavioral switches relates to the behavior of the age groups laid out in the Schulz et al. study [10], including a new capability not covered.

The subdivisions in Figure 2-2 suggests a new theoretical possibility, which I label as the "three-year-old savant" and which the Schulz paper does not yet account for. This would be a child who would be unable to integrate statistical evidence with domain-related knowledge, but would be able to deal with cross-domain relations. This theory predicts a new possibility which future psychological studies could investigate.

# Chapter 3

# Genesis System Background

One of the core ideas of this thesis is that storytelling and story understanding are qualities that are uniquely human, so in order to mimic human intelligence using computers, stories are a good place to start. The Genesis Story Understanding System is a computational model of human story understanding capabilities and the core software platform for this thesis. It has been developed by Patrick Winston's research group as a platform to advance artificial intelligence through story understanding. The system makes use of Boris Katz's START natural language parser [6] in order to read short text-based stories consisting of simple english sentences.

Once Genesis has read and parsed the story, it builds up an elaboration graph. This graph is a visual representation of the events and causal or logical connection (both explicit and implicit) that happen in the story (see Figure 3-1 below). The explicit events are the things that are explicitly stated in the text of the story while implicit events are the things that Genesis concludes about the explicit events based upon the rules that it has. For example, in Figure 3-1 there is an explicit element "Macbeth murders Duncan". This is an event that was stated in the text of the story Macbeth. Connected to that box is an implicit event "Duncan begins to be dead" which is not mentioned in the text of the story. This happens because the Genesis system has a rule that says that murdering a person causes that person to become dead.

In addition to simply displaying and matching these rules, Genesis currently has

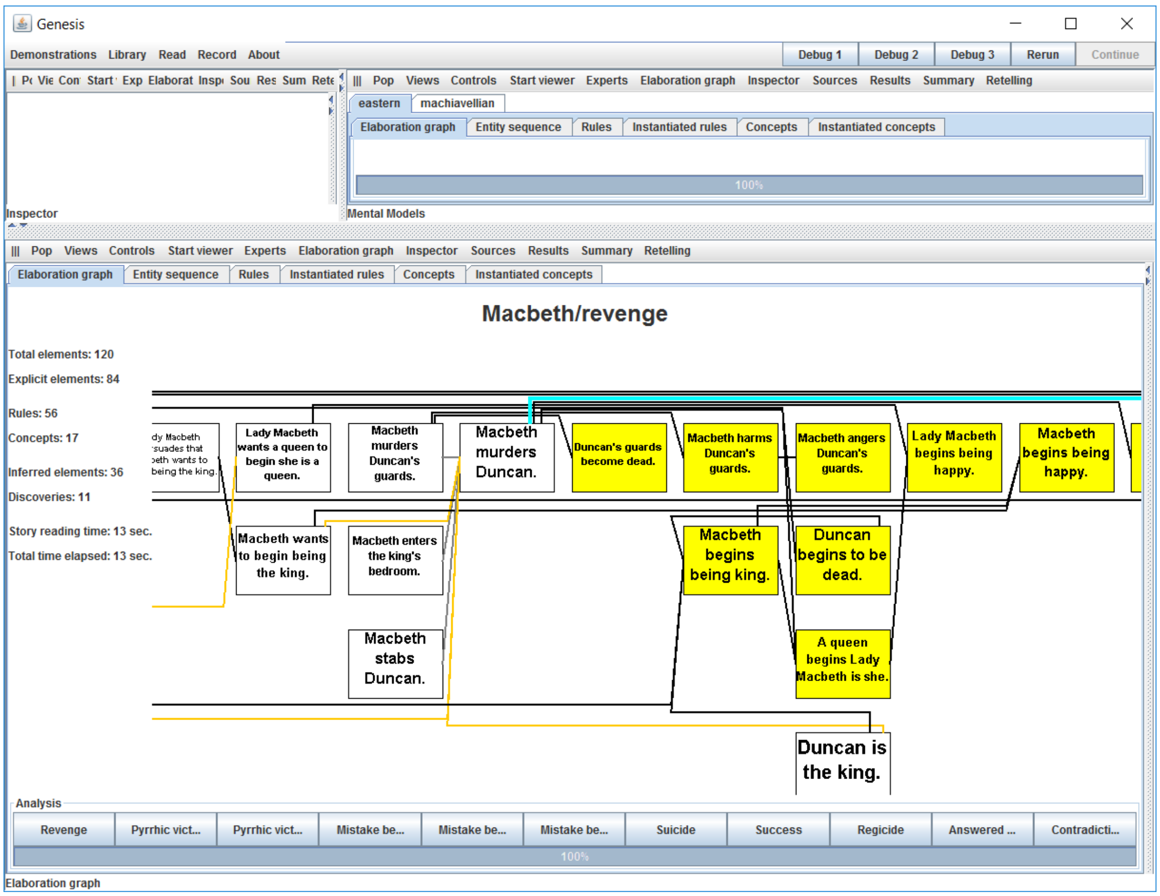

Figure 3-1: Elaboration graph displaying both explicit and implicit events from a simple version of the story Macbeth. The white boxes denote explicit events while the yellow boxes denote implicit events, and the lines show causal or logical connections between the events. For example, “Macbeth murders Duncan” causes “Duncan begins to be dead” to happen which leads to “Macbeth begins being king” etc.

the ability to understand and interpret stories in different ways. For example, it can read a story from an "eastern" perspective and from a "western" perspective and come to different conclusions as to why various events happened as a result. Similarly, Genesis can read stories through the lens of allegiance. This means that, when cued by a human user, Genesis can read a story about a conflict or a war from the perspective of either side and reach different interpretations. For example, when reading a story about a cyber-conflict between Russia and Estonia, Genesis changes its interpretation between "aggression of a bully" or "teaching a lesson" when looking at the story from the Estonian and Russian perspectives respectively. Finally, Genesis can also be persuasive in its storytelling. It can take an existing story and alter it to make characters look better or worse to the reader.

What my thesis adds to this suite of capabilities is the ability to answer causal reasoning questions in ways that model developing children.

In order to create my model, I leveraged two specific aspects of the Genesis system: rules and concept patterns.

## 3.1 The Rules

In order for the Genesis System to understand stories beyond the explicit events, we can provide knowledge to Genesis in the form of various rules. While Genesis supports numerous types of rules, the two types I used extensively in my thesis were prediction rules and explanations rules.

### 3.1.1 Prediction Rules

Prediction rules introduce implicit effects or results of events into the system's understanding of a story. For example, we reflexively infer that killing someone means that that person is now dead. However, a computer system would not know this intrinsically, so the way we supply that knowledge is through prediction rules.

In Figure 3-2, I provide an example prediction rule. When this rule is included in the reading of a story, the system will try to match the antecedent with an event in

```
XX and YY are persons.
If XX kills YY then YY becomes dead.
```

Figure 3-2: Prediction rules expressed in simple English produce implicit knowledge and connections into the story. Example borrowed from [13].

the story. In this case that means if an event like “Macbeth kills Duncan” appears in the story, then this rule will be activated, and the event “Duncan becomes dead” would be inserted into the story. Not only would “Duncan becomes dead” be inserted, but a causal link between the two events would appear in the elaboration graph [13]. Figure 3-3 shows an example of how this looks in the elaboration graph.

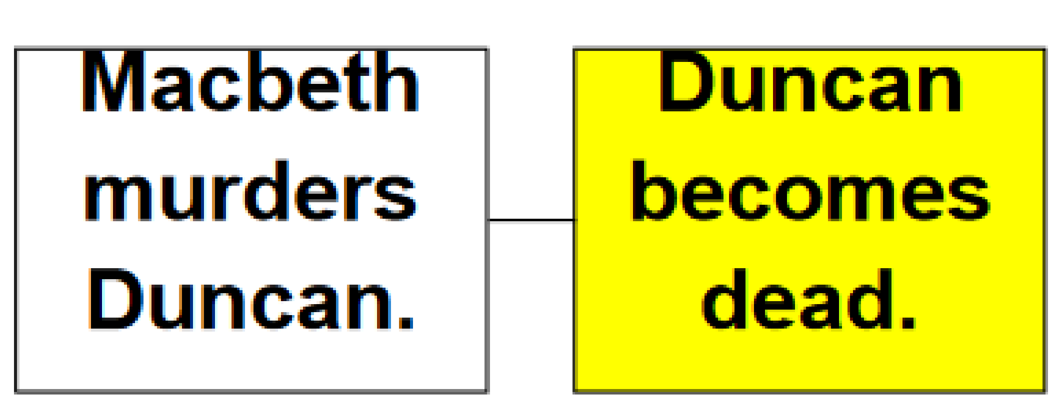


Figure 3-3: In the elaboration graph, white boxes denote explicit events and yellow boxes denote implicit events. Prediction rules supply the implicit connections and events. Example borrowed from [13].

### 3.1.2 Explanation Rules

Explanations rules enable Genesis to guess at potential causes of events (both explicit and implicit) in a story. For example, we know that angering someone might possibly lead to that person retaliating against whatever angered them. Explanation rules are the way we provide such knowledge.

While the example in Figure 3-4 may include an extreme form of retaliation, the principle still stands. When this kind of rule is included in the reading of a story, the system will search for the existence of both the antecedent and consequent within the

```
XX and YY are persons.
If XX angers YY, then YY may kill XX.
```

Figure 3-4: Explanation rules expressed in simple English suggests implicit causes for unexplained events. Example borrowed from [13].

story's events. If, and only if, the consequent has no pre-existing explanation, then the system will draw a potential causal connection between these two events. Figure 3-5 shows an example of how this looks in the elaboration graph.

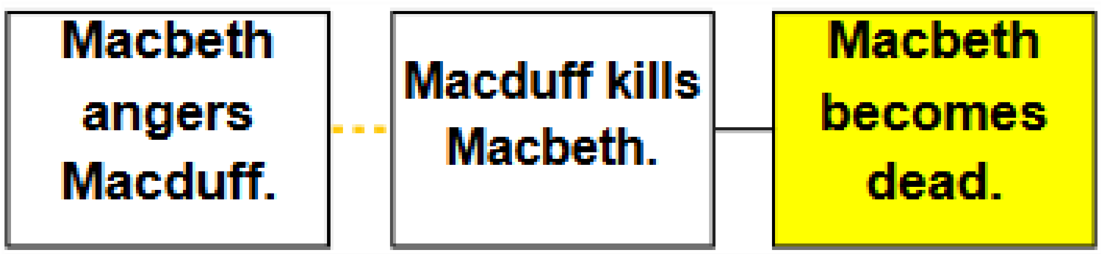


Figure 3-5: In the elaboration graph, white boxes denote explicit events and yellow boxes denote implicit events. An Explanation rule, shown here with a dashed orange line, finds a potential explanation for "Macduff kills Macbeth". Example borrowed from [13].

## 3.2 Concept Patterns

In addition to generating causal connections between individual events, Genesis can recognize concept patterns over the course of events in order to fully understand a story. For example, Genesis can find the concept of revenge occurring in the story of Macbeth even though the word "revenge" never appears.Concepts like revenge take place over the course of multiple events and can usually be summarized as one event leading to another event further down the road. Genesis uses concept patterns to identify such high-level themes in stories.

Figure 3-6 contains an example of how we might define the concept of revenge. It doesn't matter how many intermediate events there are: as long as there is some causal chain stretching from the antecedent to the consequent, then Genesis recognizes

```
XX and YY are persons.
Start description of ‘‘Revenge’’.
XX’s harming YY leads to YY’s harming XX.
The end.
```

Figure 3-6: Genesis’s concept pattern for revenge links harm to reciprocal harm through any number of intermediate events. Example borrowed from [13].

an instance of revenge. Figure 3-7 shows how a concept like revenge is realized in the elaboration graph.

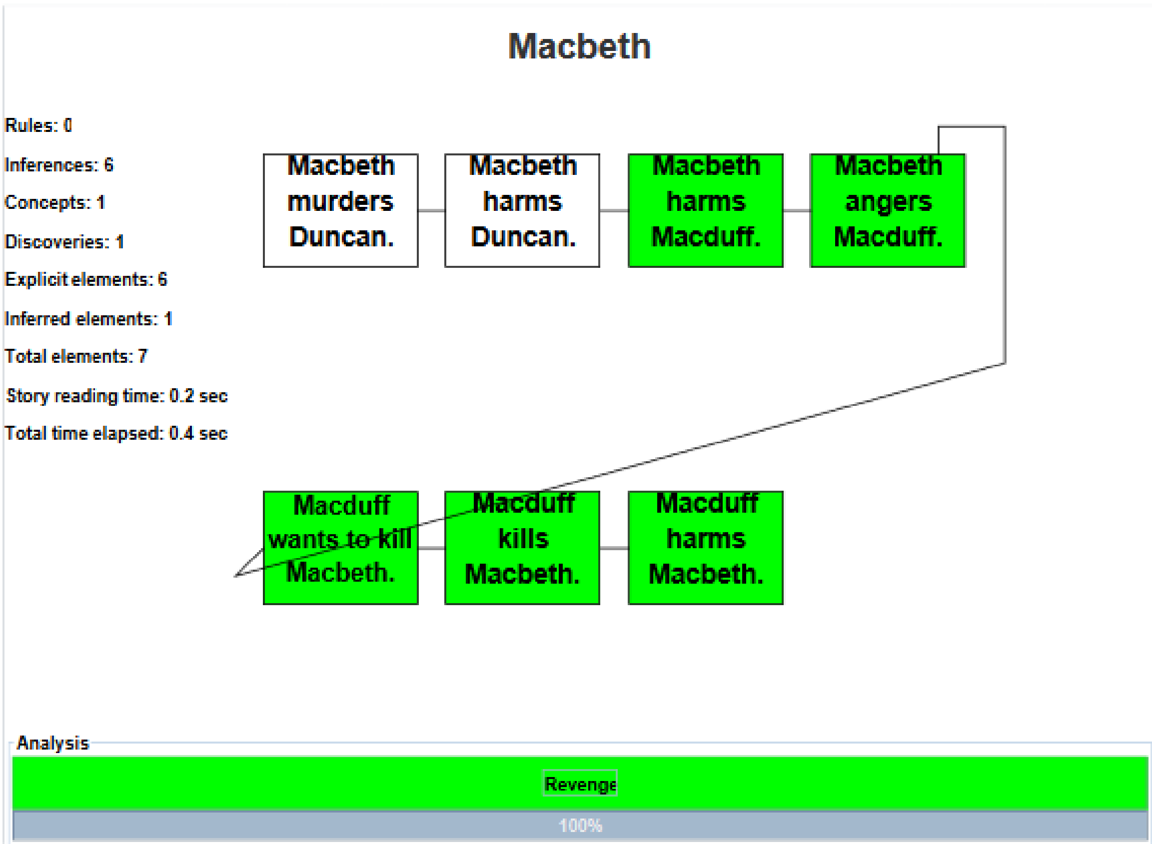


Figure 3-7: An instance of the revenge concept pattern, found in the story Macbeth, is shown in green in the elaboration graph. Example borrowed from [13].

Using Genesis’s framework of rules and concept patterns, I implemented my model described in this thesis.

# Chapter 4

# Building The Computational Models

In order to build the computational characterization of the children's behavior, I broke the implementation down into three major steps. In this chapter I discuss these steps: implementing a behavior level system based on the Bayesian model presented in the Schulz et al. study, implementing a psychologically plausible (explanation level) model, and how I made the methodology from the Schulz et al. study Genesis compatible.

My main objective in this thesis is to develop a computational, humanly-plausible account for how domain-specific causal reasoning develops in children. To start, I simply replicate the Bayesian behavior-level model by Schulz et al. [10]. I then articulate a number of more psychologically plausible explanations: that younger children have difficulty integrating different types of knowledge, that older simply have greater working memory, that censor rules initially inhibit cross-domain causes, that children reorganize their knowledge categories over time, or that knowledge in distinct brain regions becomes increasingly interconnected based on spatial proximity. I ultimately choose to implement this last, neuronally inspired, theory. Besides being especially predictive and plausible, the neuronally inspired theory has the unique virtue of being easily expressible in terms of story understanding capabilities such as rule matching and concept pattern identification.

## 4.1 The Behavior-Level Model

As a starting point for my implementation work, I decided to implement a behavior-level model to demonstrate base-line computational viability. In the original study, the authors provide a description of the Bayesian Model that they made to predict the children's behavior on the causal reasoning tasks [10]. I began by taking the mathematical description of that model [10] and implementing it on top of the Genesis system. The model reads in a story or a sequence of stories and answers causal questions such as "Did X happen because of Y or Z?". Additionally, although the original model in the paper was only intended to function for the four-to-five-year-old category of children, so I extended the statistical reasoning model to also work accurately for the three-year-old and three-and-a-half-year-old categories as well.

### 4.1.1 Three-year-old

According to Schulz et al. three-year-olds don't incorporate statistical evidence and can't make cross-domain connections when reasoning about stories. For example, a three-year-old does not believe that feeling scared can cause a stomach-ache no matter how many times those event co-occur in stories. When asked to choose between two plausible within-domain candidates, three-year-olds seem to ignore any statistical evidence from the story.

Schulz's Bayesian model, and my implementation of it, supposes that because children do not systematically use statistical evidence, each individual child effectively chooses an answer at random. As a result, this model was implemented based upon random chance and probability. While this predicts the aggregate behavior, I believe that it does not adequately or appropriately model the possible mechanisms that might be at play in an individual child's mind, which is why I developed the explanation level model described in section 4.2.

I implemented the behavior of the three-year-old model in a method called *answerAs3YearOld* that can, once Genesis has read in a sequence of stories, answer questions of the form described in section 4.3.2, meaning questions of the form: "Did

Z because X or Y?. Within this question structure, there are three cases that this model can handle: X is in the same domain as Z but Y is not, Y is in the same domain as Z but X is not, and both X and Y are in the same domain as Z. This presumes that all three events in question occur in the relevant story. If only one of X and Y happens in the story, the model will choose the one that occurs as a reasonable, though experimentally unverified, default option. Assuming that all of the events exist in the story, this model only needs to know whether the causes in question (X and Y) are within-domain or cross-domain with respect to the effect (Z) in order to answer the way a three-year-old would. The prior probabilities of choosing a cross-domain or within-domain cause were taken from the Bayesian Model given the the Schulz et al. study [10]. Specifically, the authors proposed an *a priori* bias against cross-domain causes of 0.1, and the resulting data from the three-year-old participants in the study aligns with this bias, meaning that I can produce comparable result by randomly answering the cross-domain answer 10% of the time in the three-year-old model.

This model is able to replicate the results found in the Schulz et al. study as showing in Figures 4-1 and 4-2.

```
>>>  Question: Did Your fun bambi that has them because it runs in the garden has itchy spots because it
 runs in the cattails., Answer: Your fun bambi runs in the cattails. : expert.CarolinesExpert
 .answerAs3YearOld(CarolinesExpert.java:257)
```

```
>>>  3 tally 610 : expert.CarolinesExpert.questionDummy(CarolinesExpert.java:168)
```

Figure 4-1: This is the console output from running the three-year-old behavior level model on the within-domain story. At the top, you can see the question asked and the answer given by this instance (representing one child in the study). The bottom shows the cumulative tally of responses across this and the previous 15 responses. This means that the model produced 6 "garden" responses and 10 "cattails" responses where cattails is the "correct" response.

```
>>>  Question: Did Bunny that has the bunny's tummy_ache because bunny feels scared has the bunny's 
tummy_ache because bunny eats a carrot., Answer: Bunny eats a carrot. : expert.CarolinesExpert
.answerAs3YearOld(CarolinesExpert.java:286)
```

```
>>>  3 tally 214 : expert.CarolinesExpert.questionDummy(CarolinesExpert.java:168)
```

Figure 4-2: This is the console output from running the three-year-old behavior level model on the cross-domain story. At the top, you can see the question asked and the answer given by this instance (representing one child in the study). The bottom shows the cumulative tally of responses across this and the previous 15 responses. This means that the model produced 2 "feels scared" responses and 14 "eats a carrot" responses where feeling scared is the "correct" response.

### 4.1.2 Four-to-five-year-old

The four-to-five-year-olds in the study managed to integrate statistical evidence into their causal reasoning process in both within-domain and cross-domain circumstances. Their ability to do so made them the most complex age group to model. Fortunately, the Bayesian Model in the Schulz et al. study provided the exact mathematical description of the behavior I wanted to model. As a result, I implemented an event tabulating program which counts the occurrences and co-occurrences of events in order to make Bayesian inferences.

The Schulz et al. Bayesian model revolves around the equation [10]:

$$P(h|D) = \frac{P(D|h) * P(h)}{\Sigma_{H'} P(D|h') * P(h')}$$

where $h$ is a hypothesis about the causal process, $H$ is the space of all hypotheses, $D$ is the data. In the equations below, $A$ is the potential cause being asked about [10]. Furthermore, the study looks into the specific case of a forced choice between two options. In this case, the Bayesian Model makes use of the following equation [10]:

$$P(chooseExplanationA|D) = \frac{P(ExplanationA|D)}{P(ExplanationA|D) + P(ExplanationB|D)}$$

This equation gives a numerical answer to how likely it will be that a child would choose explanation A as opposed to explanation B as their belief given data D. In this case,

$$P(ExplanationA|D) = \Sigma_H P(ExplanationA|h) * P(h|D)$$

which defines the likelihood of Explanation A being chosen given the data [10].

My program calculated the relevant probabilities by counting events in the story. Specifically, as my expert reads a story, it keeps track of the different events that have occurred in the story, the count of how many times each event has happened in the story, and finally what events happen in individual scenes. I make use of this information to calculate all the parts of the equations above. In particular, the number of times that a given event co-occurs in a scene with the effect in question is important to calculate the probability of the data given a causal hypothesis. I based my prior probability of a causal hypothesis upon the values given to the Bayesian Model in the Schulz et al. study. Specifically, I scale the probability of a cross-domain cause by 0.1, meaning that I set the $P(h)$ multiplier as $P(h) = 0.1$ when $h$ is a cross-domain cause, as suggested in the study [10]. This effectively penalizes cross-domain theories in a way that mimics conservative learning behaviors.

Ultimately these equations combined with the data from reading the stories would calculate the likelihood that a four-to-five-year-old would answer the first option from the question posed. For example in the question "Does Z because X or Y?", the model would produce a number for how likely it is that the model would answer "X". I then used this number along with a random number generator to determine whether the model answers X or Y.

Using this model, run 16 times to simulate that 16 children of this age group in the study, I was able to produce similar results to those in the study.

```
>>>  probChooseA:  0.9192938209331651 : expert.CarolinesExpert.answerAs4YearOld(CarolinesExpert.java:581)
>>>  Question: Did Your fun bambi that has them because it runs in the garden has itchy spots because it
 runs in the cattails., Answer: Your fun bambi runs in the cattails. : expert.CarolinesExpert
 .answerAs4YearOld(CarolinesExpert.java:589)


>>>  4 tally 16 : expert.CarolinesExpert.questionDummy(CarolinesExpert.java:170)
```

Figure 4-3: This is the console output from running the four-to-five-year-old behavior level model. At the top, you can see the probability that the model will choose "cattails" (the correct answer) as well as the question asked and the answer given by this instance (representing one child in the study). The bottom shows the cumulative tally of responses across this and the previous 15 responses. This means that the model produced 0 "garden" responses and 16 "cattails" responses where cattails is the "correct" response.

```
>>>  probChooseA:  0.4413793103448277 : expert.CarolinesExpert.answerAs4YearOld(CarolinesExpert.java:581)
>>>  Question: Did Bunny that has the bunny's tummy_ache because bunny feels scared has the bunny's
 tummy_ache because bunny eats a carrot., Answer: Bunny eats a carrot. : expert.CarolinesExpert
 .answerAs4YearOld(CarolinesExpert.java:589)


>>>  4 tally 709 : expert.CarolinesExpert.questionDummy(CarolinesExpert.java:170)
```

Figure 4-4: This is the console output from running the four-to-five-year-old behavior level model on the cross-domain story. At the top, you can see the probability that the model will choose "eats a carrot" (the wrong answer) as well as the question asked and the answer given by this instance (representing one child in the study). The bottom shows the cumulative tally of responses across this and the previous 15 responses. This means that the model produced 7 "feels scared" responses and 9 "eats a carrot" responses where feeling scared is the "correct" response.

### 4.1.3 Three-and-a-half-year-old

The three-and-a-half-year-olds in the study were able to integrate statistical information into their causal reasoning process when all of the options were in the same domain, but were not able to do so when there was cross-domain reasoning involved. Essentially, this meant that I needed to handle within-domain causes systematically differently than cross-domain causes. As a result, my approach to this model was a mixture of the implementation of the three-year-old model and the four-to-five-year-old model.

When answering a question, my three-and-a-half-year-old model checks whether the cause and effect are in the same domain or not. If they are, the model will integrate the statistical information. Otherwise, it will ignore the statistical evidence and make a decision based upon its prior probability of a cross-domain cause. With this implementation, the model produces similar results to that of the three-and-a-half-year-old age group in the study.

```
>>>  3.5 probA 0.875 : expert.CarolinesExpert.answerAs3AndHalfYearOld(CarolinesExpert.java:402)
>>>  Question: Did Your fun bambi that has them because it runs in the garden has itchy spots because it
 runs in the cattails., Answer: Your fun bambi runs in the cattails. : expert.CarolinesExpert
 .answerAs3AndHalfYearOld(CarolinesExpert.java:407)



>>>  3.5 tally 214 : expert.CarolinesExpert.questionDummy(CarolinesExpert.java:169)
```

Figure 4-5: This is the console output from running the three-and-a-half-year-old behavior level model on the within-domain story. At the top, you can see the probability that the model will choose "cattails" (the correct answer) as well as the question asked and the answer given by this instance (representing one child in the study). The bottom shows the cumulative tally of responses across this and the previous 15 responses. This means that the model produced 1 "garden" responses and 15 "cattails" responses where cattails is the "correct" response.

```
>>>  Question: Did Bunny that has the bunny's tummy_ache because bunny feels scared has the bunny's
 tummy_ache because bunny eats a carrot., Answer: Bunny eats a carrot. : expert.CarolinesExpert
 .answerAs3AndHalfYearOld(CarolinesExpert.java:515)



>>>  3.5 tally 115 : expert.CarolinesExpert.questionDummy(CarolinesExpert.java:169)
```

Figure 4-6: This is the console output from running the three-year-old behavior level model on the cross-domain story. At the top, you can see the question asked and the answer given by this instance (representing one child in the study). The bottom shows the cumulative tally of responses across this and the previous 15 responses. This means that the model produced 1 "feels scared" responses and 15 "eats a carrot" responses where feeling scared is the "correct" response.

## 4.2 The Psychologically Plausible (Explanation-Level) Model

While statistical reasoning may certainly play a part in our reasoning and decision making, as discussed previously, it is my fundamental belief that we are more than just statistical machines. It is my belief that any model must also be explainable, and this section of my thesis is dedicated to how I built this explainable model.

### 4.2.1 Discussion of Psychologically Plausible Models Considered

Over the course of the thesis, my primary goal has been to create models that are psychologically plausible, meaning that they could serve as a computational theory of what might actually be happening in the children's minds when they are answering the questions in the experiment. The probabilistic models described in section 4.1 only functioned as a behavioral model, not the psychologically plausible model that I set out to achieve. Those early models served as a proof of concept and a jumping off point for what I believe to be the most interesting part of the thesis: coming up with different computational theories about what might be happening during the two stages of causal reasoning development shown in the Schulz et al. study [10].

In this subsection, I propose and evaluate five explanatory computational models for development, and I argue for their psychologically plausibility. The first two models account only for the ability to incorporate statistical evidence in within-domain cases, the third accounts only for the ability to reason about cross-domain causes, and the last two account for both.

1. New Ability to Integrate Domain-Specific Information (Three to Three-and-a-Half-Year-Old Transition)

    - Knowledge Integration Theory

This theory is based upon the idea that young children may have difficulty integrating different types of knowledge in order to make decisions. For example, there was an experiment in which it was shown that young children cannot integrate color and geometric knowledge in order to make an appropriate decision [5]. Here, I propose that something similar might be happening with regard to numerical or frequency data integrating with domain knowledge.

Essentially, the way that the data from the Schulz et al. experiment can be interpreted is that the change in causal reasoning behavior between ages three and three-and-a-half is due to an increased ability to integrate numerical data across within-domain cause and effect pairs. The three-year-olds do not appear to take any of the frequency data into account when answering the questions about the cause of and effect. They effectively answer at chance for the within-domain case despite the overwhelming evidence for one option over the other. However, the three-and-a-half year olds do seem to make the connection between the frequency information and their answer to the questions.

With this theory, I propose that the development that occurs in this 6-month period can be described as a new ability to handle multiple types of data at the same time, rather than only being able to use one type of information or the other to make a decision. Although I do not pursue this theory, I note that this theory could be modeled in the Genesis system by first introducing a method of both gathering and inhibiting the use of numerical data in the three-year-old model. Then, to model development, the method of inhibition would be systematically removed to make the model behave like a three-and-a-half-year-old.

One weakness of this theory is that it cannot account for three-and-a-half-year-olds' inability to integrate statistical evidence with cross-domain causal reasoning. That issue is not a prohibitive fault of this theory, but it would be preferable for a model to account for both transitions. A second weakness of this theory is the inhibitor switch mechanism. It implies that there is some time during development where the behavior and information processing capabilities suddenly and immediately change which I find implausible.

- Working Memory Theory

  In this case, I am hypothesizing that perhaps three-year-old children have a more limited working memory than three-and-a-half-year-olds. According to this theory, three-year-olds can remember or count how many times one event or another happens, but do not have enough working memory to count how many times multiple events happen. In contrast, the three-and-a-half-year-olds have enough memory capacity to be able to keep track of the multiple pieces of data required to make within-domain statistical inferences. This working memory increase could be due to an increase in myelination [2] or because of an increase in brain size.

  In layman's terms, this is a theory is suggesting that the main developmental improvement between age three and three-and-a-half is the ability to work with more data simultaneously. While this theory explains why three-and-a-half-year-olds can integrate statistical evidence in within-domain cases, it does little to shed light on why this would only affect within-domain learning rather than both within-domain and cross-domain. In fact, if I were to choose to adhere to this theory, I would then be reasonably compelled to answer that very question: why would a memory improvement like this only correspond to within-domain causal relations? In truth, I cannot answer that question without there being some other underlying mechanism or explanation beyond a simple memory space increase. Therefore, I rejected this theory.

2. New Ability to Handle Cross-Domain Information (Three-and-a-Half to Four-to-Five-Year-Old Transition)

- Removal of a Censor Rule

  The premise of this theory is that in younger children there is some sort of “censoring rule or mechanism that inhibits or prevents children from considering cross-domain explanations. There could be many mechanisms underlying censor rules, such as under-connected neural pathways between areas of domain knowledge [9]. Essentially the lack of connectivity could be enforced in order to build up distinct domain knowledge first without the potential errors that could arise from less constrained learning at the beginning.

  The primary weakness of this theory as I see it is that it implies a more or less “instantaneous switch between behaviors. While children sometimes exhibit rapid qualitative shifts in behavior [1], it does not seem to me that causal reasoning behaviors would fall under this pattern. I think that it is most likely a slow, continuous development.

3. The Ability to Integrate Statistical Evidence and Cross-Domain Information (Three-Year-Old Through Four-to-Five-Year-Old Transitions)

   - Neuronally Inspired Theory

     According to this theory, domain-specific knowledge and processes lie in isolated regions of the brain. Conceptual connections form initially within domains, and then later between domains, as a direct result of increased neural connectivity.

     Conceptually, you have“clusters of domain knowledge that are stored separately from each other, and as you grow, these clusters become more connected both internally and between each other. By way of explanation, consider two circles that each contain some dots. Each of the circles represents a domain and the dots are all little bits of knowledge or examples of things that belong in the domain circle that they are within. I

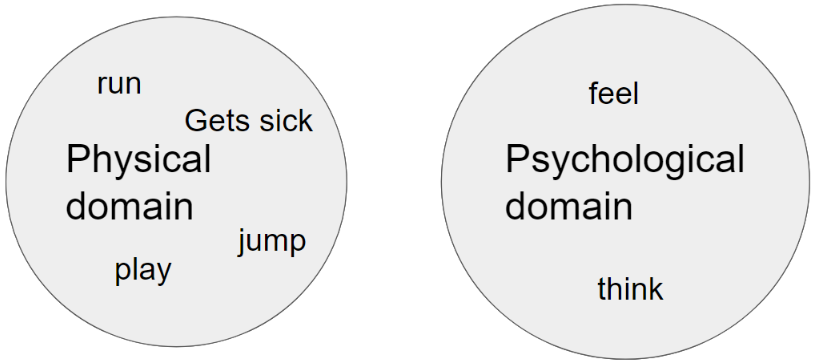


Figure 4-7: Three-year-olds' domain-specific knowledge can be represented by distinct regions with distinct pieces of conceptual knowledge within them. This abstract representation explains three-year-olds' lack of ability to integrate statistical evidence within domains or to recognize abstract connections between physical events and psychological events.

propose that children progress through three states with two transition or developmental periods, as follows.

The starting state consists of physically disconnected domains with few connections within domains and no connections between domains. In the context of the Schulz et al. [10] experiment, one circle contains physical domain knowledge and the other contains psychological domain knowledge 4-7.

In the case of my Genesis based behavior level models, the domain knowledge contained would be the different verbs that belong to each domain. This initially disconnected model accounts for the question answering capabilities of a three-year-old child. Specifically, because there are no"connections between the dots within the circles, they don't have the ability to link or combine data about different dots within the same circle. This would provide an explanation as to why they don't seem to use the evidence provided in the story to do better than guessing at chance for within-domain cause-effect pairs. Additionally, this model would give an explanation as to why there is such a heavy bias towards within-domain versus cross-domain explanations at three-years-old; because the cause and effect in a cross-domain pair are completely separate from each other in the mind of the child. The cause would be in one circle while the effect was in the other.

In this second state, connections have been formed within domains due to proximity (Figure 4-8). In terms of the model for this theory, this development would mean that new connections between dots within the same circle have been formed. The connections between concepts here could convey many sorts of information including how the concepts generally relate to each other or even how often you notice the concepts appearing together. Therefore, the first transition in this model would be the growth or appearance of these new within-domain connections that allow us to use the evidence provided in the stories for within-domain cause-effect pairs, while still penalizing cross-domain cause-effect pairs. This state captures how three-and-a-half-year-olds can begin to incorporate statistical information within domains, but not between domains.

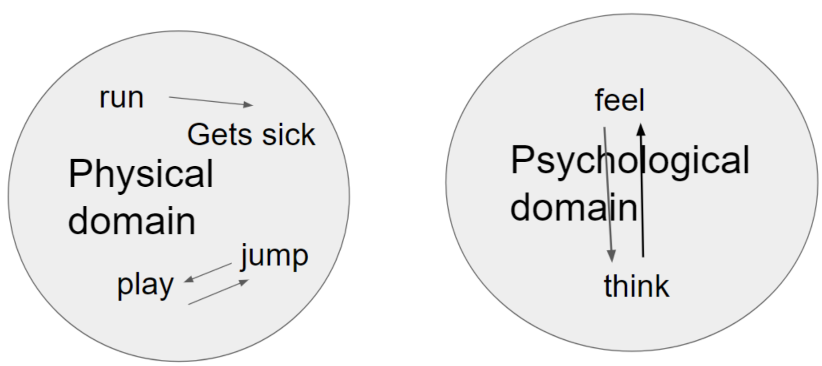


Figure 4-8: Three-and-a-half-year-olds' domain-specific knowledge can be represented by distinct regions with connected pieces of conceptual knowledge within them. This abstract representation explains three-and-a-half-year-olds' ability to integrate statistical evidence within domains and their lack of ability to recognize abstract connections between events across the two domains.

Finally, in the third state, connections have also begun to form between spatially distinct domains. Essentially, between three-and-a-half and four-to-five-years-old I propose that new connections are forged between the

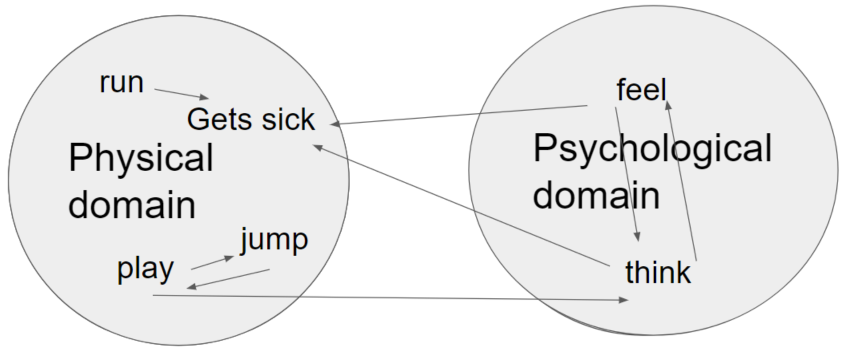


Figure 4-9: Four-to-five-year-olds' domain-specific knowledge can be represented by distinct regions with connected pieces of conceptual knowledge within and between them. This abstract representation explains four-to-five-year-olds' ability to both integrate statistical evidence within domains and to recognize abstract connections between events across the two domains.

dots in the different circles meaning that concepts from the physical domain become connected and related to concepts in the psychological domain (Figure 4-9). The fact that this new kind of cross-domain connection now exists is what enables four-to-five-year-olds to be able to consider a physical action to be the cause of a psychological effect. Looking at this through the lens of this theory, the four-year-olds' new ability to learn cross-domain cause-effect relations would signify that new connections were formed between the different circles.

All in all, I find this theory and neuronal analogy to be very compelling. That said, it is not without its drawbacks. For example, in reality, there must be some way to curate, prune, and/or prevent different connections from forming. The limiting of connections between ideas is important because it prevents children from learning irrelevant or incorrect information [9]. However, the theory and approach that I have laid out here does nothing to address this important aspect of learning.

Even so, this is a good first step. It may not be as predictive as we might want it to be, as it doesn't predict which information falls into which domains, which domains are formed and when, and which connections are formed and when, but it gives a framework that enables us to think and talk about domain based causal reasoning development.

- Search Through Probability Space of Different Models Rather than Causes

  This theory suggests that what is happening during the development stages is a change in the probability space of mechanisms rather than a change in the mechanism itself. To elaborate, in the initial stage of development that I discuss in section 4.1, I had written the models as various probabilistic searches through the solution space. The various models weight options differently based upon the "age" of the model to mimic the behavior of the various age groups of children from the study. What I am suggesting with this theory is that perhaps it isn't just the probability of individual causes that change, but rather the whole conceptual framework or mechanism that they are thinking about.

  I propose here that a plausible explanation of what is happening during development is that the conceptual framework changes, not just the probability of individual causes. Essentially, I am saying that perhaps there are a few different mechanisms for understanding causal relations that all exist during all of these ages, however due to the child's beliefs, some mechanisms are preferred over others. Therefore, there are some methods of solution that have"higher weights in terms of probability than others, and these weights are what change over time. This theory would explain the overall shift in answers across age groups, but it could also explain the outlier responses in each case as simply a probabilistic error rather than a systematic failure.

  While this theory is interesting, I find it problematic because it does nothing to improve the actual mechanisms of the models. It only really discusses the transitions rather than the states themselves which leaves too many loose ends to be a reliable theory. For this reason, I have decided not to pursue or implement this theory.

### 4.2.2 Representational Choice

I decided to implement the Neuronally Inspired Theory which has the virtues of being explanatory, compositional, and sensible. It is also easily implementable using explanation and prediction rules and concept patterns. I supplied Genesis with domain knowledge such as that if you play with an object, then you contact it and that contact is a physical concept.

Earlier, in the behavior-level model (section 4.1), I had hard coded the necessary domain knowledge by categorizing verbs such as “run” or “think” as physical or psychological. I was able to drastically improve upon this representation in the explanation-level model. By using rules and concept patterns I was able to impart more interesting and relevant domain based knowledge to my model such that it would be more explainable and thus psychologically plausible. These rules and concept patterns essentially enable my model and Genesis to represent the formation of new connections between concepts or events. Therefore, they may abstractly represent new neuronal connections and thus enable or prevent the ability to reason well about cross-domain information.

### 4.2.3 Rules for the Model

One of the important aspects of Genesis that I leveraged for this thesis is the ability to define personality traits. When writing stories for Genesis, you can designate whether or not Genesis should be reading the story through the lens of specific personalities. The personality trait files can define things like additional or different rules for how to interpret events that happen in the main story. This was exactly the functionality that I wanted to have in order to implement the models for the different age groups through rules. I created “personality trait” (for this use case, I think of them more as “ability level”) files for each age group: three-year-old, three-and-a-half-year-old, and four-to-five-year-old.

Initially, I had set out to design a set of rules for each age group, however there was such significant overlap between them that I created a fourth file containing what

I considered to be common age knowledge that all of the individual models would need. The common age knowledge encompasses information about causal relations within domains. In particular, it defines the knowledge required to make a complex logical inference such as that running through cattails can cause itchy spots: children must know that running is a type of physical action, that running in something can cause you to contact it, that contact is also a physical action, and that contact can cause itchy spots. Figure 4-10 shows an example of some of the prediction rules in the common age knowledge file.

```
VV is a thing.
XX is a person.

// PHYSICAL ACTION --> CONTACT
If XX experiences a physical-action in VV, then XX contacts VV.
If XX experiences a physical-action with VV, then XX contacts VV.

// ACTION --> PHYSICAL-ACTION
If XX runs in VV, then XX experiences a physical-action in VV.
If XX plays with VV, then XX experiences a physical-action with VV.

// FOR BUNNY STORY
If XX eats VV, then XX experiences a physical-action with VV.
```

Figure 4-10: Examples of Prediction Rules Used in Explanation Level Models

In addition to the prediction rules defined above that will always create causal links, I also found a need for explanation rules. When certain events occur they *can* be a cause for further events, but are not necessarily direct indicators that some effect *must* occur as a result. In the case of this study, just because Bambi comes into contact with something does not inherently mean that Bambi will get itchy spots. However, the models should be able to identify that contacting something could be the cause for the itchy spots. This is why the explanation rules are necessary. Figure 4-11 shows an example of some of the explanation rules in the common age knowledge file.

The rules in the common age knowledge file, enable my computational model to

```
VV is a thing.
XX is a person.

// DEFINING SICKNESSES

Sickness is a kind of physical-thing.
Sickness is a kind of psych-thing.

Physical-sickness is a kind of sickness.
Psych-sickness is a kind of sickness.
Itchy spots is a kind of physical-sickness.
Nausea is a kind of physical-sickness.

If XX contacts VV, then XX may experience a physical-sickness.

//DEFINING EMOTION

Emotion is a kind of psych-thing.
Negative-emotion is a kind of emotion.
Positive-emotion is a kind of emotion.
Scared is a kind of negative-emotion.
Great is a kind of positive-emotion.

CC is an emotion.

If XX experiences a psych-action about VV, then XX may feel CC.
```

Figure 4-11: Examples of Explanation Rules in Explanation Level Models

produce causal links from the potential causes in the stories to the effect in question in the stories. It turned out that there was only one rule that needed to be different in the four-to-five-year-old model from the other two (which are identical). Specifically, the four-to-five-year-old model needed to add a link between psychological experiences and physical ones. In order to do so, I added the explanation rule shown in Figure 4-12 to the four-to-five-year-old model.

With all of these rules in place, my model of the different age groups produces elaboration graphs like the ones shown in Figure 4-13 and Figure 4-14. These figures show the elaboration graphs for the three-year-old and three-and-a-half-year-old mod-

```
VV is a thing.
WW is a person.
If WW experiences a negative-psych-action,
  then WW may experience physical-sickness.
```

Figure 4-12: Examples of Explanation Rules in Four-to-Five-Year-Old Explanation Level Model

els and the four-to-five-year-old model respectively. The primary difference between the children is that the four-to-five-year-olds infer that feeling scared (a psychological event) can cause nausea (a physical event), while younger children do not. The two elaboration graphs shown in Figures 4-13 and 4-14 show how my model captures this difference through the presence or absence of an explanation rule. This one connection is what enables the "older" model to understand that a psychological action may be able to produce a physical effect.

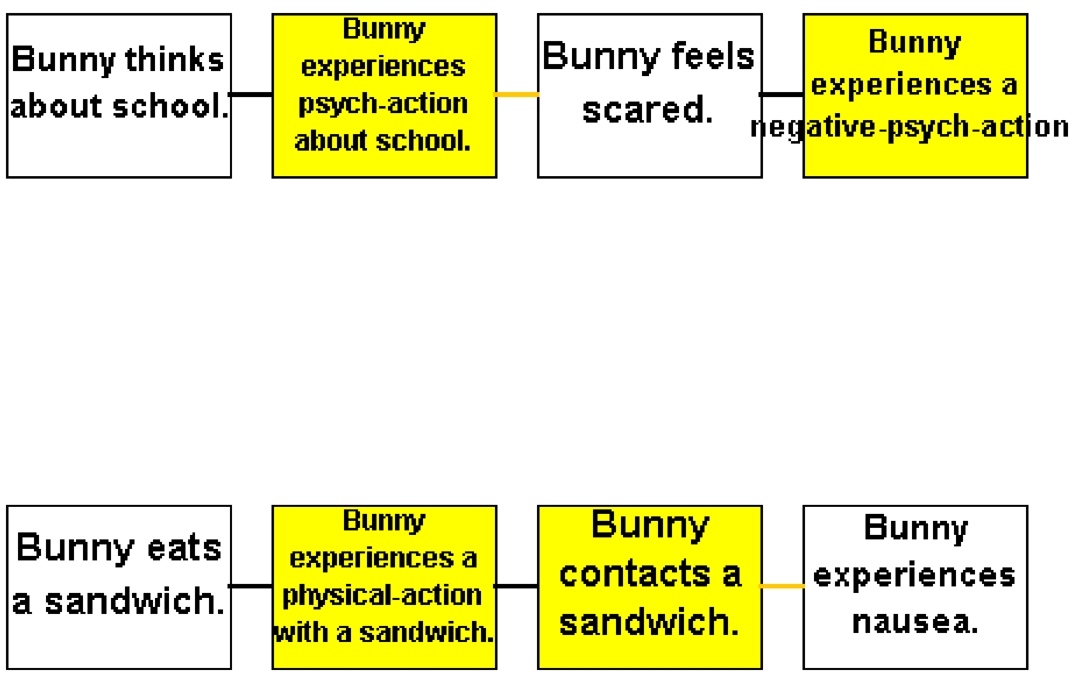


Figure 4-13: When presented with a sequence of stories both three and three-and-a-half-year-olds are limited to within-domain explanations: there are no connections between the top (psychological) events and the bottom (physical) set of events although they are part of the same story.

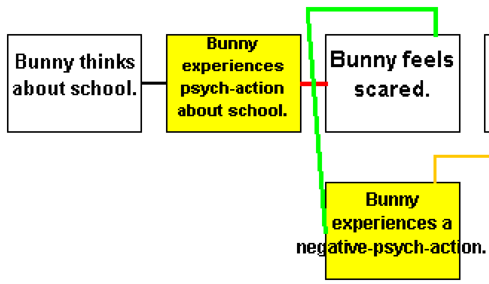


Figure 4-14: In contrast to Figure 4-13, a four-to-five-year-old who reads the same story concludes that being scared (a psychological event) can cause nausea (a physical event). In my system, this new possibility is achieved by introduction of a cross-domain explanation rule.

### 4.2.4 Concept Patterns

Despite the chain of causal connections that clearly exist in the elaboration graphs in Figures 4-13 and 4-14, this is not enough for the Genesis system to make that causal deduction. In order to have Genesis recognize a causal chain, concept patterns must be employed. To this end, I defined a few different concept patterns to recognize different causal links. Specifically, I defined patterns to identify where running in something is linked to a physical-sickness, where eating something is linked to a physical-sickness, and where feeling something is linked to a physical-sickness (Figure 4-15).

With these concept patterns, Genesis recognizes when physical or psychological causes eventually lead to physical or psychological effects. These connections are consequently highlighted in the elaboration graph that the model produces. Importantly, the psychological cause to physical effect concept pattern only gets recognized in the four-to-five-year-old model's understanding of the story as shown in Figure 4-17 as opposed to Figure 4-16.

Furthermore, the concept patterns were important to create for the explanation level models because their existence enabled the underlying statistical calculation to be more reasonable as well. Now, instead of simply relying upon whether or not two events happened in the same scene, I can set the calculations to keep track of

```
XX is a person.
VV is a thing.

Start description of ‘‘Physical cause ingestion to physical effect’’.
XX’s eating VV leads to XX’s experiencing a physical-sickness.
The end.

Start description of ‘‘Physical cause contact in to physical effect’’.
XX’s running in VV leads to XX’s experiencing a physical-sickness.
The end.

Start description of ‘‘Psychological cause to physical effect’’.
XX’s feeling VV leads to XX’s experiencing a physical-sickness.
The end.
```

Figure 4-15: Concept Patterns for Sickness

how many scenes activated relevant concept patterns. This means that the statistical reasoning can now be restricted to candidate causes rather than all prior events which I think is a much improved method of reasoning.

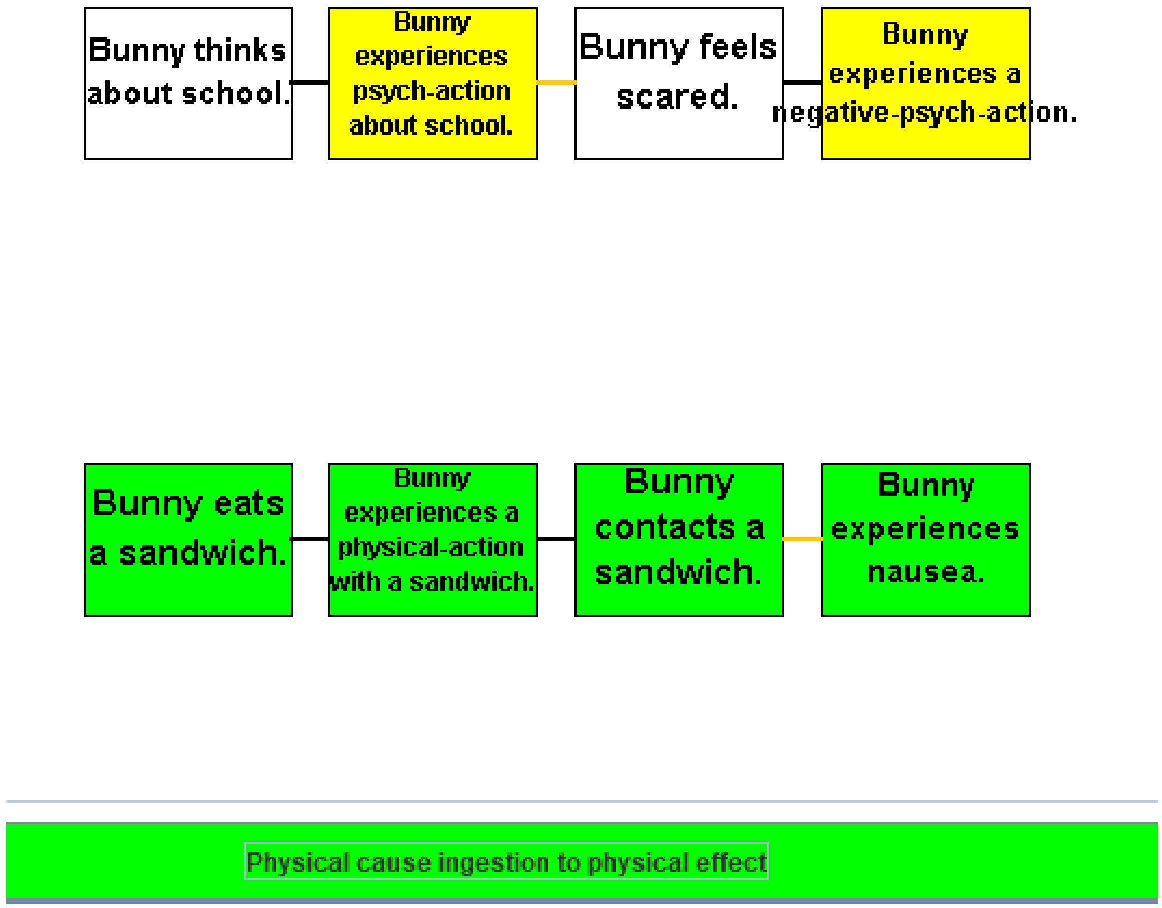


Figure 4-16: The physical cause to physical effect concept pattern highlights a long distance connection between eating a sandwich and experiencing nausea.

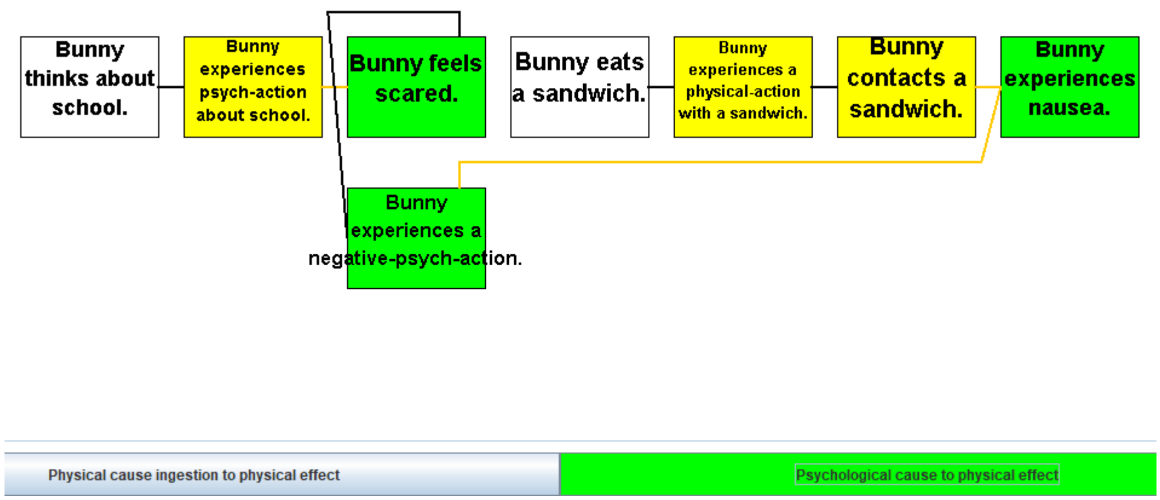


Figure 4-17: The four-to-five-year-old model has a psychological cause to physical effect concept pattern which highlights a long distant causal connection between feeling scared and experiencing nausea.

## 4.3 Converting the experiment to be Genesis Compatible

In order to develop the models that you have just seen, the first crucial task was to adapt the experiment to the Genesis Story Understanding System. At its most basic, the Schulz et al. study consisted of reading some stories and then answering questions about them, at which point, the Genesis system would need to be able to read the stories from the study and answer the kinds of questions asked in the study. While Genesis technically had the ability to read stories and answer questions before my work, both the stories and questions from the study needed to be tweaked to work with the Genesis system. The simplicity of children's stories belies the wide variety of skills required to understand them.

### 4.3.1 The stories

The stories included in the Schulz et al. study are very simple from a human perspective, however there are a decent number of aspects of the stories that are complicated for Genesis to handle. In particular, I ran into some problems with the START parser, some issues with the repetitive nature of the story, and some difficulties with how Genesis was dealing with matching events and actors. Some of the problems ended up being addressed by altering the story file itself while others wound up finding solutions in the actual Java code or uncovering underlying issues in the Genesis system that needed to be worked around while a fix for the larger bug was in progress.

On the whole, the stories were fairly easy to recreate as the Schulz et al. study provided a full text of the stories in its appendices [10]. The vast majority of the stories were able to parse well after simply copying the text of the stories from the study. However, there were a few specific parts of the stories that wouldn't parse, so I changed them in as minimal a way as possible that enabled them to parse. For example, the phrase "cedar trees" would not parse properly, so I changed it to "pine trees" which parsed fine. Some other minor changes like this included:

- ... runs **through** a place → ... runs **in** the place.
- ... plays with **his toy** truck → ... plays with **a** truck.
- ! → .
- show and tell → show-and-tell
- tummy ache → tummy-ache → nausea
- splitting compound sentences into two separate sentences. For example: "On Thursday afternoon, Bambi plays jump rope and Bambi runs in the sand." → "On Thursday afternoon, Bambi plays jump rope. Bambi runs in the sand."
- "On Thursday afternoon, Bambi plays jump rope." → "The time is Thursday afternoon. Bambi plays jump rope."
- removing "on his legs" from "Bambi has itchy spots on his legs."

Genesis does not have a full fledged, human-level capacity to reason about time. The stories from the study all rely on the reader being able to draw a distinction between what happened on one morning or afternoon versus the next. Genesis does not have the ability to do this without adding a specific scene change idiom into the story. Specifically, starting a sentence with "Then," informs Genesis that a new scene is starting [13]. The addition of scene markers into Genesis enables it to understand that repeated events are independent events, not the same event simply being mentioned twice. This was an essential functionality to have due to the repetitive nature of the stories. As a result, I added these scene markers to the story at the beginning of each morning and afternoon which produced the following kind of change in sentences:

"The time is Thursday afternoon." → "Then, the time is Thursday afternoon."

A full text of the stories can be found in Appendix A.

### 4.3.2 The questions

Much like with the stories, there were some minor alterations that I needed to make to the questions. The original question format in the study was:

"Why does Bambi have itchy spots? Is it because of running through the garden or because of running through the cattails?"

This kind of format, unfortunately does not work in the Genesis system, seeing as it is phrased as two separate questions. In order to make this question work for the Genesis system, I rephrased it into a single question:

"Does bambi have itchy spots because bambi runs in the garden or bambi runs through the cattails?"

Initially, I had tried repeating the "because" to more clearly identify that it's inquiring about two separate causes, however the START parser was unable to handle questions of that format. The same structure of question was used for both stories, so the same sort of alteration to the question for the second story was made.

Having described how I tailored the original study to be parsable by Genesis, I now describe the architecture of my models and how they are able to understand the inputs.

## 4.4 High-Level Structure

The first step was to create a workspace that was integrated with the Genesis System. The Genesis System was designed in such a way as to make modular development easy and easy to set up. Essentially, there is a box-and-wire design, so information is transferred between modules via "wires" where you can specify sources and destinations [12]. I made use of this design in order to construct my models as their own module within the Genesis System.

### 4.4.1 Expert Class and Wiring

I structured my module as an expert module within Genesis. What this means is that in the main Genesis GUI, there is a checkbox that indicates whether or not this code is being included when reading stories. Due to the way that I wired up the expert, the story information will be received and handled only if the expert's checkbox is checked. This setup enables my models to be run by anyone with the Genesis system, as opposed to needing to run from a specific file. Implementing the models in this manner turned out to be particularly convenient during the transition between developing the behavioral level model and developing the explanation level model. I was able to simply "plug in" the new model file rather than needing to wire up a whole new setup.

### 4.4.2 Filtering and Answering Questions

I implemented question answering in such a way that questions would only be transmitted to my model if a corresponding radio button is selected in the Genesis GUI. I wanted this feature to be included in the system so that I could be certain that the questions would be directly transmitted to my expert class for processing rather than Genesis's general question answering system. When questions are sent to my expert class, I do some pre-processing on the questions to make sure that they are in the relevant format for the study. Specifically, I implemented a filter in my expert class such that it will only attempt to answer questions that are well-formed for the study. This means that the system is designed to only answer questions of the form "Does ZZ because XX or YY?". Essentially, I needed to first check to make that that the question is a "Did" question, and then check if there is a two element disjunction within the question.

Once I have determined that the question being asked is one that is relevant to the experiment, I pass the question along to my model so that it can answer the question.

# Chapter 5

# Contributions

In this thesis I have proposed that causal reasoning is an essential aspect of human intelligence and that we can understand it by studying how it develops. I have proposed that story understanding capabilities play an essential role in how we humans make causal connections, integrate information from past experience, and organize information according to domain. To advance these ideas, I built a humanly-plausible model of how causal reasoning develops in distinct stages, replicating a study by Schulz et al. [10] and taking it to another level. Unlike many existing cognitive models, my model not only *describes* the developmental phenomenon, but *explains* it in terms of story-understanding mechanisms. Furthermore, it exhibits the three essential traits of psychologically plausible models—it is explanatory, compositional, and sensible.

More specifically, I:

- Introduced three criteria that a computational model must satisfy in order to be *psychologically plausible*. These criteria and this thesis can serve as a model for future endeavors into computationally modeling psychological phenomena.

- Implemented a naive, purely statistical, *behavior-level* model of developmental causal reasoning, recreating the experimental setup and results of the Schulz et al. experiment [10].

- Argued that the behavior-level model accurately describes the behavior of the children, but does explain the behavior or how it develops.

- Defined and evaluated various explanatory models, presenting the pros, cons, and psychological plausibility arguments for each.
- Proposed an *explanation-level* cognitive theory which explains how children's cross-domain casual reasoning ability develops over time. I expressed this theory in terms of ability to integrate statistical information and ability to include cross-domain cause-and-effect.
- Abstracted the explanation-level theory into a switch-like mechanism and showed how it predicts a novel type of causal reasoning behavior.
- Implemented a psychologically plausible computational model of the explanation-level theory using a story understanding approach that leverages the Genesis system's rule and concept pattern matching abilities.
- Demonstrated that the model adequately reproduces the results from the Schulz et al. study [10] and takes it to another level by explaining the behavior in terms of psychologically plausible story understanding mechanisms such as rule inference and concept pattern matching.

# Appendix A

# Stories

```
Start story titled "Bambi Story".

Bambi likes to prance and run in lots of different places.
Running is fun for Bambi.
//Then on Monday morning, Bambi runs in the pine grove.
Then, the time is Monday morning.
Bambi runs in the pine grove.
Bambi gets excited.
Bambi runs in the cattails.
Bambi experiences itchy spots. // on his legs.  removed for effect matching purposes
Bambi does nothing. // placeholder hack for mental model 2 not transmitting the last

//Then on Monday afternoon, Bambi runs in the pine trees. //new line cedar trees
Start story titled "afternoon".
Then, the time is Monday afternoon.
Bambi does not experience itchy spots.
Bambi runs in the pine trees.
Bambi plays on the rope swing.
Bambi feels great.
Bambi does not experience any itchy spots.
```

```
Bambi does nothing.

//Then on Tuesday morning, Bambi gets excited.
Start story titled "Tuesday morning".
Then, the time is Tuesday morning.
Bambi gets excited.
Bambi runs in the cattails.
Bambi runs in the grass.
Bambi experiences itchy spots. // on his legs.  removed for effect matching purposes
Bambi does nothing.

//Then on Tuesday afternoon, Bambi reads a book. //new line - and
Start story titled "Tuesday afternoon".
Then, the time is Tuesday afternoon.
Bambi reads a book.
Bambi runs in the rock_bed. //through rock bed
Bambi feels great.
Bambi does not experience any itchy spots.
Bambi does nothing.

//Then on Wednesday morning, Bambi runs in the marsh.
Start story titled "Wednesday morning".
Then, the time is Wednesday morning.
Bambi runs in the marsh.
Bambi gets excited.
Bambi runs in the cattails.
Bambi experiences itchy spots. // on his legs.  removed for effect matching purposes
Bambi does nothing.

//Then on Wednesday afternoon, Bambi runs in the apple orchard. //new line - and
```

```
Start story titled "Wednesday afternoon".
Then, the time is Wednesday afternoon.
Bambi runs in the apple orchard.
Bambi plays with a truck. //through apple... with his toy truck
Bambi feels great.
Bambi does not experience any itchy spots.
Bambi does nothing.

//Then on Thursday morning, Bambi gets excited.
Start story titled "Tursday morning".
Then, the time is Thursday morning.
Bambi gets excited.
Bambi runs in the cattails.
Bambi runs in the leaves.
Bambi experiences itchy spots. // on his legs.  removed for effect matching purposes
Bambi does nothing.

//Then on Thursday afternoon, Bambi plays jump rope. //new line - and
Start story titled "Thursday afternoon".
Then, the time is Thursday afternoon.
Bambi plays jump rope.
Bambi runs in the sand.
Bambi feels great.
Bambi does not experience any itchy spots.
Bambi does nothing.

//Then on Friday morning, Bambi runs in the bushes.
Start story titled "Friday morning".
Then, the time is Friday morning.
Bambi runs in the bushes.
```

```
Bambi gets excited.
Bambi runs in the cattails.
Bambi experiences itchy spots. // on his legs.  removed for effect matching purposes
Bambi does nothing.

//Then on Saturday morning, Bambi gets excited.
Start story titled "Saturday morning".
Then, the time is Saturday morning.
Bambi gets excited.
Bambi runs in the cattails.
Bambi runs in the grass.
//Bambi gets excited.
//Bambi runs in the cattails.
Bambi experiences itchy spots. // on his legs.  removed for effect matching purposes
Bambi does nothing.

//Then on Saturday afternoon, Bambi gets his hair brushed. //new line - and
Start story titled "Saturday afternoon".
Then, the time is Saturday afternoon.
Bambi gets his hair brushed.
Bambi runs in the blueberry_patch. //through blueberry patch
Bambi feels great.
Bambi does not experience any itchy spots.
Bambi does nothing.

//Then on Sunday morning, Bambi runs in the garden. //through garden
Start story titled "Sunday morning".
Then, the time is Sunday morning.
Bambi runs in the garden.
Bambi gets excited.
```

```
Bambi runs in the cattails.
Bambi experiences itchy spots. // on his legs.  removed for effect matching purposes
Bambi does nothing.

//Then on Friday afternoon, Bambi runs in the playground. //new line - and
Start story titled "Friday afternoon".
Then, the time is Friday afternoon.
Bambi runs in the playground.
Bambi roller skates. //through playground
Bambi feels great.
Bambi does not experience any itchy spots.
Bambi does nothing.



Then, on the next day Bambi's spots were all gone.
Have fun Bambi.

The end.

//Does bambi have itchy spots because bambi runs in the cattails or bambi runs in the
Start story titled "Bunny Story".

Bunny is a rabbit. Bunny is a girl.

// SHOW-AND-TELL --> SCHOOL FOR MATCHING PURPOSES
// This is Bunny.
Bunny is scared because next week she has to give school.
school makes Bunny scared.

//Then on Monday morning, Bunny thinks about school.
```

```
Then, it is Monday morning.
Bunny thinks about school.
Bunny feels scared.
Bunny eats some cheese.
Bunny experiences nausea.

The end.


Start story titled "Monday afternoon".
//Then on Monday afternoon, Bunny ties her shoes. // and
Then, the time is Monday afternoon.
Bunny ties her shoes.
Bunny eats strawberries.
Bunny feels great. //! can't parse !
Bunny does not experience nausea.

The end.


Start story titled "Tuesday morning".
//Then on Tuesday morning, Bunny eats a popsicle.
Then, the time is Tuesday morning. //split so can match...
Bunny eats a popsicle.
Bunny thinks about school.
Bunny feels scared.
Bunny experiences nausea.

The end.


Start story titled "Tuesday afternoon".
//Then on Tuesday afternoon, Bunny eats some toast. // and
```

Then, the time is Tuesday afternoon.

Bunny eats some toast.

Bunny takes a bath.

Bunny feels great. //! can't parse !

Bunny does not experience nausea.

The end.

Start story titled "Wednesday morning".

//Then on Wednesday morning, Bunny thinks about school.

Then, the time is Wednesday morning.

Bunny thinks about school.

Bunny feels scared.

Bunny eats French fries.

Bunny experiences nausea.

The end.

Start story titled "Wednesday afternoon".

//Then on Wednesday afternoon, Bunny plays bingo. // and

Then, the time is Wednesday afternoon.

Bunny plays bingo.

Bunny eats pasta.

Bunny feels great. //! can't parse !

Bunny does not experience nausea.

The end.

Start story titled "Thursday morning".

//Then on Thursday morning, Bunny eats a muffin.

Then, the time is Thursday morning.

Bunny eats a muffin.

Bunny thinks about school.

Bunny feels scared.

Bunny experiences nausea.

The end.

Start story titled "Thursday afternoon".

//Then on Thursday afternoon, Bunny eats some yogurt. // and

Then, the time is Thursday afternoon.

Bunny eats some yogurt.

Bunny brushes her teeth.

Bunny feels great. //! can't parse !

Bunny does not experience nausea.

The end.

Start story titled "Friday morning".

//Then on Friday morning, Bunny thinks about school.

Then, the time is Friday morning.

Bunny thinks about school.

Bunny feels scared.

Bunny eats some soup.

Bunny experiences nausea.

The end.

Start story titled "Friday afternoon".

//Then on Friday afternoon, Bunny plays on the monkey bars. // and

```
Then, the time is Friday afternoon.
Bunny plays on the monkey bars.
Bunny eats a banana.
Bunny feels great. //! can’t parse !
Bunny does not experience nausea.

The end.

Start story titled "Saturday morning".
//Then on Saturday morning, Bunny eats a carrot.
Then, the time is Saturday morning.
Bunny eats a carrot.
Bunny thinks about school.
Bunny feels scared.
Bunny experiences nausea.

The end.

Start story titled "Saturday afternoon".
//Then on Saturday afternoon, Bunny eats some tofu. // and
Then, the time is Saturday afternoon.
Bunny eats some tofu.
Bunny builds a snowman.
Bunny feels great. //! can’t parse !
Bunny does not experience nausea.

The end.

Start story titled "Sunday morning".
//Then on Sunday morning, Bunny thinks about school.
```

Then, the time is Sunday morning.

Bunny eats a sandwich.

Bunny thinks about school.

Bunny feels scared.

Bunny experiences nausea.

The End.

Start story titled "Monday morning 2".

Then the next day Bunny gave school.

She did very well and everyone clapped!

Hurray for Bunny. //! can't parse !

The End.

# Bibliography

[1] Elizabeth Baraff Bonawitz, Tessa J.P. van Schijndel, Daniel Friel, and Laura Schulz. Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, 64:215–234, 2012.

[2] G. Dehaene-Lambertz and E.S. Spelke. The infancy of the human brain. *Neuron*, 88(1):93–109, oct 2015.

[3] Clark Glymour, Choh Man Teng, and David Danks. Bayesian ptolemaic psychology. 2015.

[4] Hyowon Gweon and Laura Schulz. 16-month-olds rationally infer causes of failed actions. *Science*, 332(6037):1524, 2011.

[5] Linda Hermer-Vazquez, Elizabeth S. Spelke, and Alla S. Katsnelson. Sources of flexibility in human cognition: Dual-task studies of space and language. *Cognitive Psychology*, 39(5):3–36, aug 1999.

[6] Boris Katz. Annotating the world wide web using natural language. In *Proc. 5th RIAO Conference on Computer Assisted Information Searching on the Internet (RIAO '97)*, 1997.

[7] Rachel W. Magid, Mark Sheskin, and Laura E. Schulz. Imagination and the generation of new ideas. *Cognitive Development*, 34:99–110, 2015.

[8] Jo Marchant. A journey to the oldest cave paintings in the world ... Smithsonian magazine website: http://www.smithsonianmag.com/history/journey-oldest-cave-paintings-world-180957685/, January 2016. Accessed October 24, 2016.

[9] Marvin Minsky. Negative expertise. *International Journal of Expert Systems*, 7(1):13–19, 1994.

[10] Laura Schulz, Elizabeth Baraff Bonawitz, and Thomas L. Griffiths. Can being scared cause tummy aches? naive theories, ambiguous evidence, and preschoolers' causal inferences. *Developmental Psychology*, 43(5):1124–1139, 2007.

[11] Aaron Sloman. *The Computer Revolution in Philosophy: Philosophy Science and Models of Mind*, chapter 2.5.5. The Harvester Press Limited, online:

http://www.cs.bham.ac.uk/research/projects/cogaff/crp/ edition, 1978. Accessed May 3, 2018.

[12] Patrick H. Winston. System building using genesis's box-and-wire mechanism. http://groups.csail.mit.edu/genesis/Documentation/wire.pdf, October 2014.

[13] Patrick H. Winston. How to speak genesese. http://groups.csail.mit.edu/genesis/Documentation/Genesese.pdf, June 2017.

[14] Patrick H. Winston and Dylan Holmes. The genesis manifesto: Story understanding and human intelligence. https://courses.csail.mit.edu/6.S079/pdf/manifesto.pdf, May 2018.

[15] Patrick Henry Winston. The strong story hypothesis and the directed perception hypothesis. In *Proc. of the AAAI Fall Symposium on Advances in Cognitive Systems*. Association for the Advancement of Artificial Intelligence, November 2011.