The Moral of the Story Is...
Extract a Generalizable Lesson from a Fable Through Emotional Explanation and Alignment with a Past Experience

by

Zhaozheng Alice Jin

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 2019

Author .....
Department of Electrical Engineering and Computer Science
May 24, 2019

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

Accepted by .....
Katrina LaCurts
Chair, Master of Engineering Thesis Committee

The Moral of the Story Is...

Extract a Generalizable Lesson from a Fable Through Emotional Explanation and Alignment with a Past Experience

by

Zhaozheng Alice Jin

Submitted to the Department of Electrical Engineering and Computer Science
on May 24, 2019, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

After reading *The Tortoise and the Hare*, it may feel instinctive to conclude that the moral of the story is "slow and steady wins the race". However, research has shown that this is not so obvious to children, who tend to focus on story-specific details like napping in the middle of the race. In learning the moral, it is crucial to generalize. Otherwise, we would need a fable for every unique circumstance.

What computational process underlies our seemingly intuitive ability to extract a generalizable moral of a story? Fables play integral roles across cultures and societies. From a young age, children are read fables to instill moral values. If we are build an artificial human intelligence system, we must first answer this question.

In this thesis, I take a step toward fulfilling my vision by building MAXIM, a new module in the Genesis Story Understanding System. From a fable and a description of a past experience written in English, MAXIM extracts and generalizes the moral of the fable and explains it in English. For example, from *Rudolph the Red-Nosed Reindeer* and *The Math Aficionado*, MAXIM concludes `It's ok if you are different` because `Valuable Rudolph is different`. Notably, idiosyncrasies of the story, such as having a red nose, are not stated in the moral.

MAXIM extracts a generalizable moral by first interpreting both stories on an emotional level. By explaining the emotional states and their transitions, the system can identify the moral challenge. Then, MAXIM aligns the stories by emotional states to abstract away story-specific details.

In developing MAXIM, I have distilled four principles for extracting a generalizable moral: *Viewpoint Character Principle*, *Reversal of Fortune Principle*, *Emotional Explanation Principle*, and *Emotional Alignment Principle*. Though internalized by adults, these principle are learned, perhaps unconsciously, by children.

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

Acknowledgments

Working on this thesis was the best of times, it was the worst of times. It continued through the winter of despair and the spring of hope. It was the epoch of belief, it was the epoch of incredulity. I did it, but I couldn't have done it alone.

Thank you **mom & dad** for loving, trusting, and supporting me through my endeavors and shenanigans, even when you weren't sure what I was up to and especially when I wasn't sure what I was up to.

Thank you **Eddie** for always patiently answering my questions, whether it was years ago when I read you stories or last week when I inquired about what you learned in school. It doesn't seem like I stop learning from you.

Thank you **Amy** for never being farther than a hyper bunny sticker away, even while you were half a world away.

Thank you **Candice** for being an attentive sounding board just next door.

Thank you **Dylan Holmes** for helping me find the right words to articulate my ideas and for your optimistic perspective.

Thank you **Professor Randall Davis** for lending a helping hand during a difficult time.

Thank you **Professor Patrick Winston** for leading with courage, optimism, and compassion. I learned to think more intentionally, to communicate my visions and contributions more eloquently, and to take care of my people. Without you, my dream of building an artificial intelligence system would remain just that—a dream.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

I love stories and I love sharing my favorites with my little brother. I recall reading him fables and asking him what he learned from them. From *The Three Little Pigs*, he learned that you should build a house with bricks so a wolf can't blow it down. I distinctly remember being at a loss of words, wondering why he refuses to generalize lessons applicable to other situations, and instead focuses on idiosyncrasies of the story. I quickly found out this was not because he took a particular liking to architecture. From *The Tortoise and the Hare*, he learned that you shouldn't take a nap in the middle of a race. In disbelief, I would ask him questions until he came to the conclusion that I thought was all too obvious.

Perhaps the moral of a story is not actually obvious. In a study by Narvaez et al., 3rd and 5th graders read a story and were asked to identify a target story with similar lesson. The results showed that the children were more likely to point to a target story with similar surface content, even if it illustrated a different moral [9].

It may feel instinctive to conclude that "slow and steady wins the race" when we are familiar with the pithy saying and the story where it comes from. In abstracting away the surface features unique to the story, we are able to recognize when a situation is analogous to the race between the Tortoise who persevered despite its disadvantage and the complacent Hare. However, for a child who is reading this story for the first time, how can he look beyond the story-specific details to extract and learn a moral that can be applied to other situations?

1.1 Vision

If we are to build an artificial human intelligence system, then we must understand the computational process that underlies our seemingly intuitive ability to extract a generalizable moral of a story.

Of all the species who have walked the earth, only we—human beings—tell stories. The Strong Story Hypothesis states that "the mechanisms that enable humans to tell, to understand, and to recombine stories separate our intelligence from that of other primates" [13]. Stories play integral roles across cultures and societies. We tell stories to entertain; we tell stories to memorialize events; we tell stories to teach life lessons. From a young age, children are read fables to instill moral values. A mark of a lesson well learned is that you are able to recognize when a situation is analogous to another. If we cannot look beyond salient story-specific details, then we would need a fable to advise us what to do in every unique circumstance. In learning the moral of a story, it is crucial to generalize.

1.2 A Look Ahead

I have taken a step toward fulfilling my vision by building MAXIM. From a fable and a description of a past experience written in English, MAXIM extracts the moral of the fable and explains it in English. For example, from *Rudolph the Red-Nosed Reindeer* and *The Math Aficionado*, MAXIM concludes

`It's ok if you are different.`

because

`Valuable Rudolph is different`.

This is captured in figure 1-1.

Notably, the original story does not explicitly state that Rudolph is different. This is inferred using common sense knowledge, after observing that Rudolph is the only reindeer with a red nose. Likewise, Rudolph's difference proving to be valuable is also inferred, rather than explicitly stated. MAXIM abstracts away story-specific

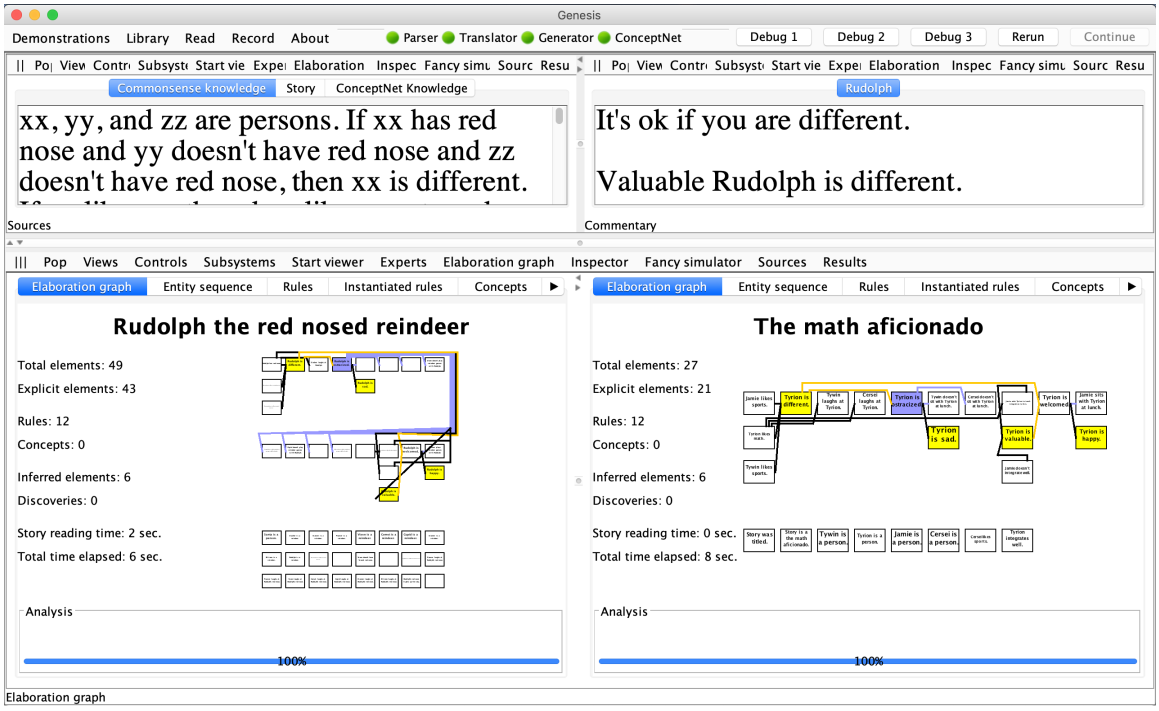

Figure 1-1: From a fable and a description of a past experience written in English, MAXIM extracts and generalizes the moral of the fable and explains it in English. The conclusion from *Rudolph the Red-Nosed Reindeer* and *The Math Aficionado* is shown in the top-right panel. The elaboration graphs for both stories are shown individually in the bottom panels. Some of the the common sense knowledge used in inferences can be seen in the top-left panel.

details such as having a red nose. As a result, the extracted moral is more general, facilitating drawing analogies to similar situations.

At a high level, MAXIM achieves this by

1. interpreting both stories on an emotional level, i.e. by inferring the characters' emotional states, such as being excluded from the reindeer game can make one sad

2. explaining each emotional state to find its initial cause, which in Rudolph's case is being the only one with the red nose

3. aligning the two stories by emotional states to find a common explanation for each emotional state, therefore abstracting away story-specific details and concluding that Rudolph being different explains why it is sad

4. grouping emotional states by character to isolate character arcs

5. identifying a character's challenge and reaction to the challenge by explaining the change in the character's emotional states

6. translating the moral and the explanation from MAXIM's inner representation to English

The implementation of each step is described in Chapter 3: MAXIM. MAXIM's implementation substrate, the Genesis Story Understanding System, is described in Appendix A: The Genesis Enterprise.

In developing MAXIM, I have distilled four principles of the computational process for extracting a generalizable moral of the story:

Viewpoint Character Principle: the moral of the story depends on the point of view.

Reversal of Fortune Principle: to find the moral of the story, look for a reversal of fortune. The change of a character's fortune for the better or worse signals the consequence of the character's choice when faced with a challenge.

Emotional Explanation Principle: to find a reversal of fortune, look for emotional arcs by explaining the emotional states and their transitions in the story.

Emotional Alignment Principle: to find the right level of abstraction for the moral, align the fable and a past experience with similar emotional arcs.

These principles are explained in detail in Chapter 2: Guiding Principles for Extracting a Generalizable Moral. They are formed by finding some commonalities among fables. Empirical grounding for these principles from psychology research are described in Appendix B: Empirical Grounding for the Guiding Principles.

MAXIM's output for other fables are showcased in Chapter 4: What Can We Learn from a Reindeer, a Shepherd Boy, a Tortoise, and a Hare?.

Chapter 2

Guiding Principles for Extracting a Generalizable Moral

Extracting a generalizable moral of a fable is difficult. It may feel intuitive when we are familiar with the fables where pithy maxims come from. However, research has shown that it is not obvious to children, who instead focus on salient, story-specific details [9]. So, what principles have adults internalized, so much so that we hardly think twice about them, but children must learn, perhaps unconsciously?

This chapter discusses four guiding principles for extracting a generalizable moral, independent of implementation. These principles are formed by finding some commonalities among fables. Empirical grounding for these principles from psychology research are explained in Appendix B: Empirical Grounding for the Guiding Principles. The key idea is to interpret a fable emotionally and compare it to a superficially different past experience with similar emotional states.

2.1 Identify Viewpoint Character

In fables, animals take on human qualities and objects inanimate in the real world come to life. Just like humans, they can think and feel. These characters are employed to illustrate lessons. In some stories, the focus is on a single character, such as Rudolph in *Rudolph the Red-Nosed Reindeer* or the Shepherd Boy in *The Boy Who Cried*

Wolf. In other stories, different lessons can be learned depending on which character's perspective a reader takes. For example, in *The Tortoise and the Hare*, the Tortoise is an epitome of slow and steady wins the race, whereas the Hare demonstrates a pitfall due to hubris. So, character arcs should be considered individually.

> **Viewpoint Character Principle:** the moral of the story depends on the point of view.

2.2 Identify Reversal of Fortune

Many fables follow a structure in which a character undergoes a challenge and the character's fortune reverses depending on the character's reaction to the challenge. In this thesis, I consider fables with a single reversal of fortune. This is based on the notion that fables are intentionally written to follow this narrative structure for the purpose of instructing a single, central moral to young children. Complex narrative structure may get in the way of learning the moral in the first few years of life. However, it is still possible for a number of morals to be embedded in a story with a complex narrative structure.

In this thesis, I also categorize a character arc as either a motivational story or a cautionary tale.

A **motivational story** is marked by a happy ending after an initial difficulty. At first, a character is presented with a difficult situation. However, the challenge turns into an opportunity, and the character's fortune changes for the better. For instance, the Tortoise from *The Tortoise and the Hare* shows us even if you are at a disadvantage, you can still persevere and emerge victorious. A motivational story provides guidance on how to navigate a difficult situation.

On the other hand, in a **cautionary tale**, a character's fortune changes for the worse. The character's initial benefit turns into a source of difficulty. As opposed to the Tortoise, the Hare shows us being complacent of your advantage can cost you. A cautionary tale warns readers of the consequence of lapse in judgement.

Reversal of Fortune Principle: to find the moral of the story, look for a reversal of fortune. The change of a character's fortune for the better or worse signals the consequence of the character's choice when faced with a challenge.

2.3 Identify Challenge by Explaining the Reversal in Emotional State

A reversal of fortune typically results in a reversal of emotional state. An upturn from initial to final emotional state indicates a **motivational story**, whereas a downturn indicates a **cautionary tale**. Explaining the cause of a character's initial emotional state reveals the challenge that the character faces. Explaining the cause of a character's final emotional state reveals the character's reaction to the challenge.

Emotional Explanation Principle: to find a reversal of fortune, look for emotional arcs by explaining the emotional states and their transitions in the story.

It should be emphasized that by explaining the reversal in a character's emotional state, a system can identify a challenge without prior knowledge of what constitutes a challenge. For example, having a red nose is not inherently a challenging characteristic: there is nothing troublesome about a red nose in and of itself. However, having a red nose is a challenge in context of *Rudolph the Red-Nosed Reindeer* because all the other reindeer do not have red noses. Rudolph's difference leads to ostracism from his peers, which in turn saddens him. By leveraging the reversal of a character's emotional state, we can identify the challenge particular to the story.

2.4 Abstract Away Story-Specific Details by Emotionally Aligning Two Stories

Emotional explanation reveals the story-specific challenge, such as having a red nose. However, this level of detail is not helpful in drawing analogies to other situations

and, thus, applying the moral. There may be a scenario where not having a red nose is actually a challenge. To extract a general moral applicable to other situations, it is not enough to examine only the fable at hand. When reading the fable, a child may recall a past, real life experience with the same change in emotional state for similar reasons. Aligning the emotional states in a fable and a description of a past experience reveals the generelizable challenge common in both stories, such as being different. Doing so, story-specific details are abstracted away, facilitating analogical reasoning.

> **Emotional Alignment Principle:** to find the right level of abstraction for the moral, align the fable and a past experience with similar emotional arcs.

Chapter 3

MAXIM

This chapter describes the implementation of MAXIM, a new module in the Genesis Story Understanding System. From a fable and a description of a past experience written in English, MAXIM extracts and generalizes the moral of the fable and explains it in English. At a high level, this is done by first inferring and explaining characters' emotional states. MAXIM emotionally aligns the fable and the past experience to abstract away story-specific details. Then, character arcs are isolated. The reversal of each character's emotional state is explained to identify the character's challenge and reaction to the challenge. Lastly, MAXIM's internal representation is translated to English, outputting the moral and explaining it with evidence from the fable.

Figure 3-1 shows the system at work, identifying moral of *Rudolph the Red-Nosed Reindeer* as `It's ok if you are different`. The explanation for this moral is `Valuable Rudolph is different`. This fable is emotionally aligned with *The Math Aficionado*, in which one student passionate about math is rebuffed from a lunch table full of sports enthusiasts. This story parallels Rudolph's exclusion from the reindeer games for uniquely having a red nose. Both stories can be found in Appendix C.1.

In the following sections, I describe each step in detail, illustrating the inner workings of the system with the help of Rudolph and the math aficionado, Tyrion. To understand the implementation of MAXIM, it is necessary to first understand its implementation substrate, Genesis. Appendix A: The Genesis Enterprise describes

the system as well as the motivation for the enterprise.

3.1 Reading a Fable and Recalling a Past Experience

To model how a reader may recall a past experience similar to the fable at hand, both the fable and the description of what may resemble a real life experience are input into Genesis's Story Processor simultaneously. Both stories are elaborated with the provided common sense knowledge and interpreted using the same process described in the following sections.

To extract the moral of *Rudolph the Red-Nosed Reindeer*, the Story Processor reads this and *The Math Aficionado*, the real life experience, simultaneously. The provided set of common sense knowledge can be found in Appendix C.1.3.

Before going into details of the moral extraction process, it is worth examining what kind of common sense knowledge is provided to the Story Processor. One of the explanation rules for the Rudolph-Tyrion pair of stories is:

```
If xx is different and xx is valuable, then xx may be welcomed.
```

Arguably, one only needs to be valuable to be welcomed, not both different *and* valuable. However, this rule is written as such to highlight that a character is valuable because he is different and to exclude other ways of being valuable, such as being talented or trustworthy. Here, I trade the generality of the common sense knowledge for the specificity of the moral. While developing a theory of common sense knowledge of value is insightful, it is beyond the scope of this work.

3.2 Inferring and Explaining Emotional States

3.2.1 Inferring Emotional States Using Common Sense Knowledge

A character's emotional state can be inferred and injected into the story by providing common sense knowledge in the form of deduction rules:

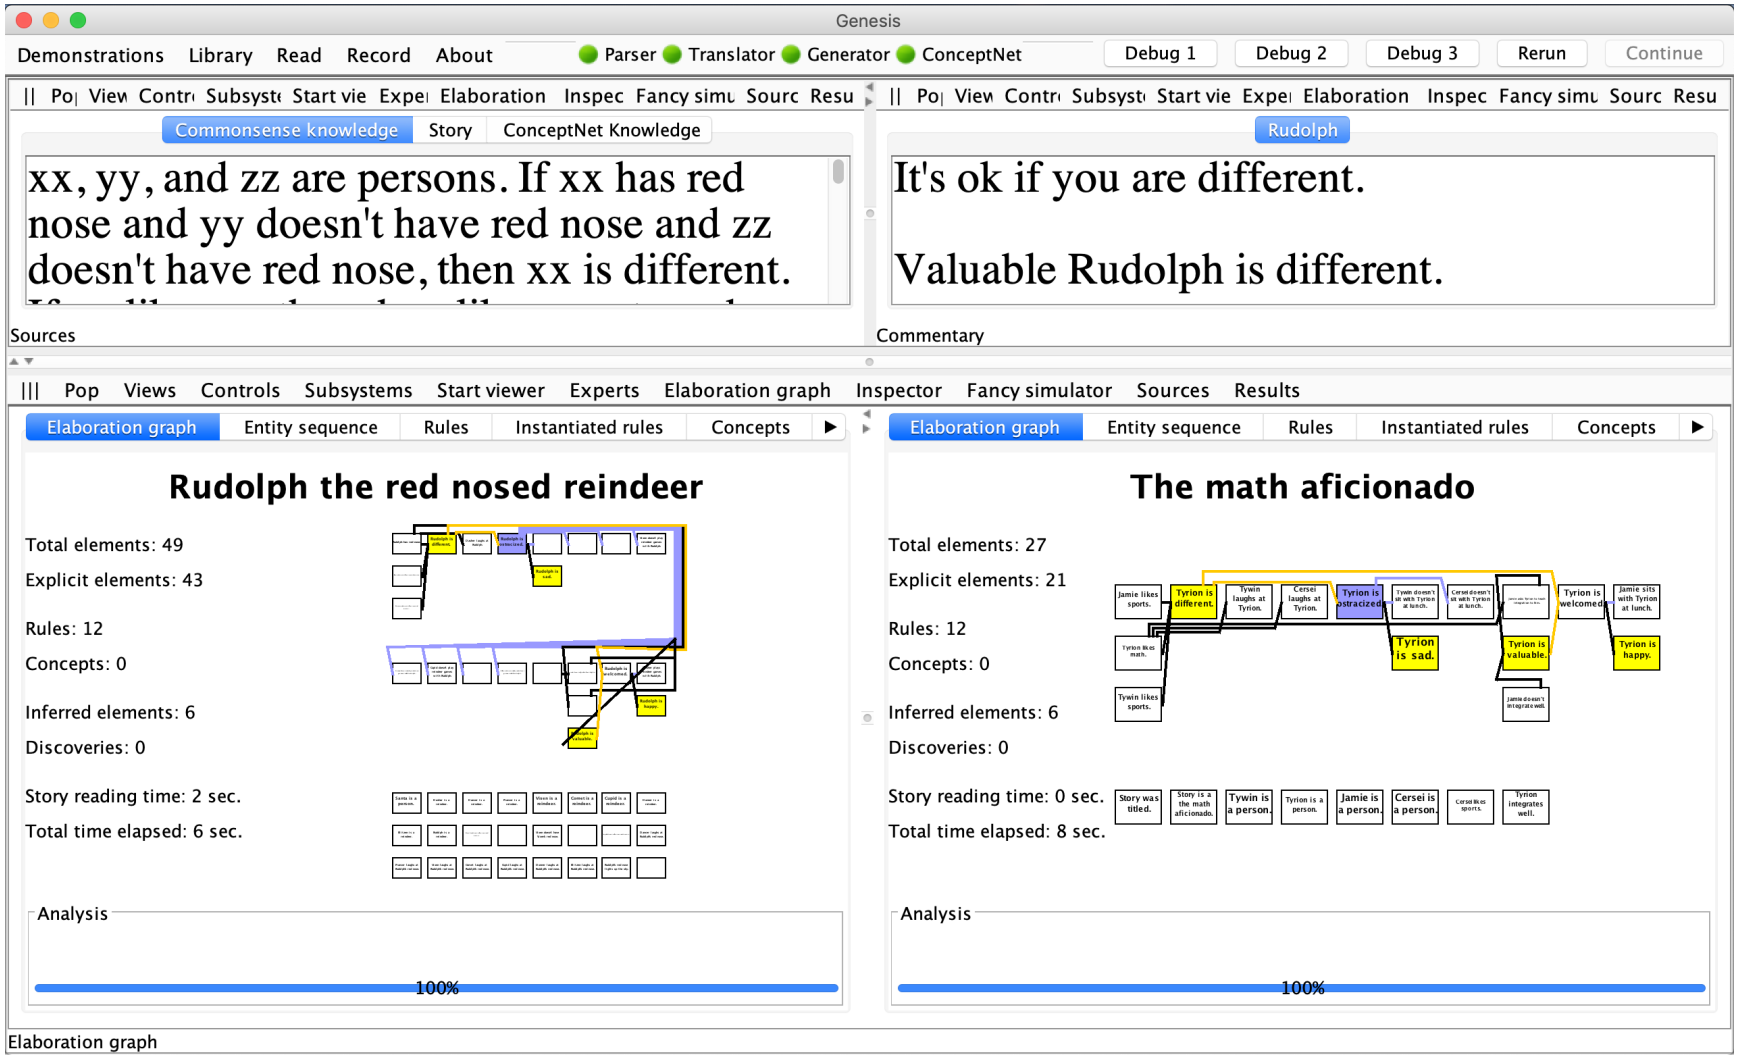

Figure 3-1: MAXIM identifies the moral of *Rudolph the Red-Nosed Reindeer* as `It's ok if you are different`, as seen in the top-right panel. This is done by aligning the fable with a description of what may resemble a real-life experience. The elaboration graphs for both stories are shown individually in the bottom panels. Both stories are explained on an emotional level, with inferences of emotional states injected into the stories. Some of the the common sense knowledge used in inferences can be seen in the top-left panel.

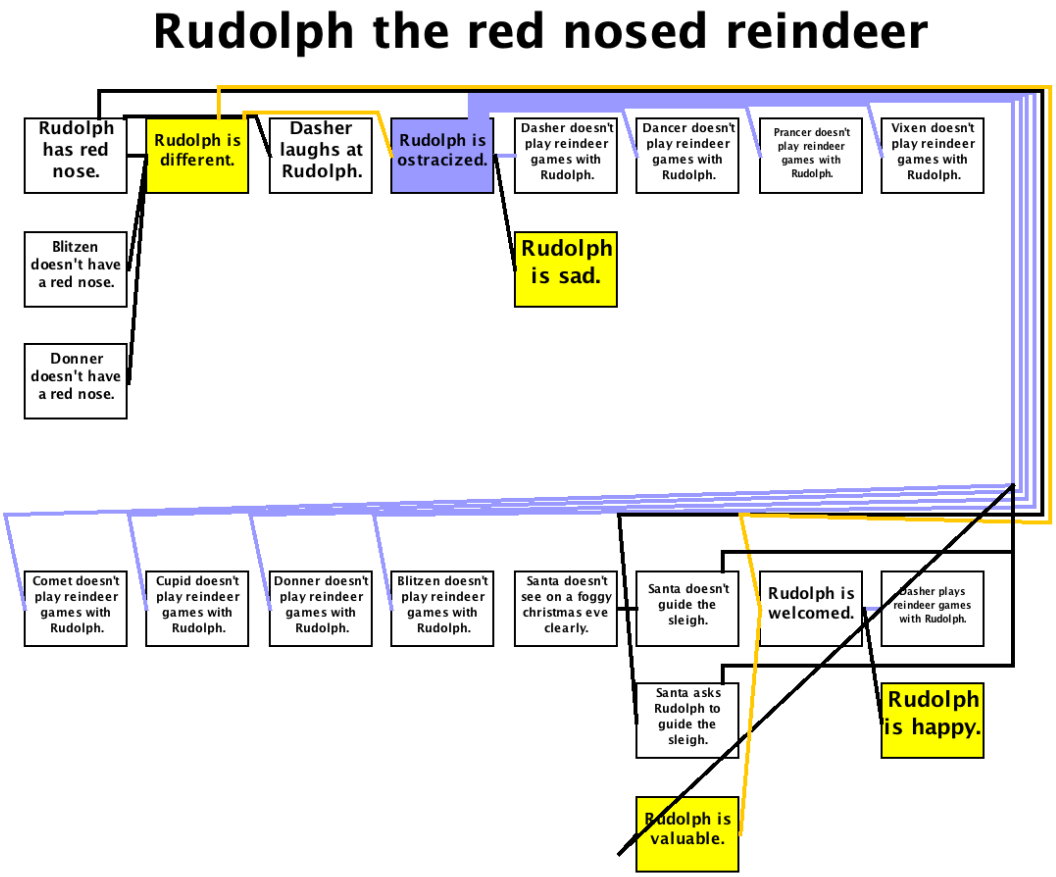

Figure 3-2: Elaboration graph for *Rudolph the Red-Nosed Reindeer*, enlarged from figure 3-1. Rudolph's emotional states—initially sad then happy—are inferred using common sense knowledge and injected into the story. These new story elements are highlighted in yellow.

```
If xx is ostracized, then xx is sad.
If yy is welcomed, then yy is happy.
```

The consequents—`xx is sad` and `yy is happy`—are matched to a character and added as story elements. The newly added elements can be seen highlighted in yellow in the story's elaboration graph, shown in figures 3-2 and 3-3.

3.2.2 Explaining Emotional States Using Backward Chaining

Following the *Emotional Explanation Principle*, MAXIM builds an emotional arc, which shows the series of events that led up to an inferred emotional state. For each

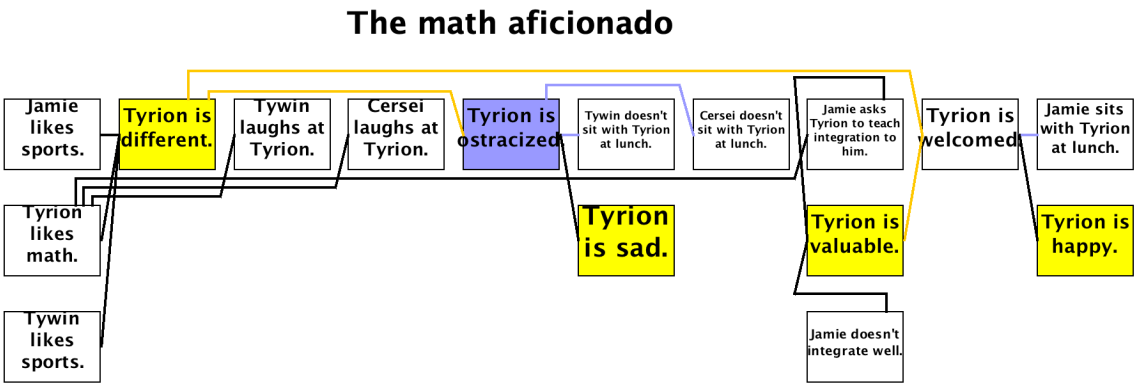

Figure 3-3: Elaboration graph for *The Math Aficionado*, enlarged from figure 3-1. Tyrion's emotional states—initially sad then happy—are inferred using common sense knowledge and injected into the story. These new story elements are highlighted in yellow.

emotional state in the story, MAXIM finds the initial cause by backward chaining from the emotional state to build the longest chain of inferences, incrementally adding story elements to the front of the chain. The emotional arc for Rudolph's sadness is shown in figure 3-4.

3.2.3 Finding a Common Explanation by Aligning Stories

In building an emotional arc, MAXIM finds the initial cause for an emotional state. However, this cause is an idiosyncrasy of the story. Following the *Emotional Alignment Principle*. MAXIM aligns the emotional arcs in the fable and the past experience to abstract away story-specific details. Starting from the endpoints of the emotional arcs, the system backtracks simultaneously in both arcs until the explanation differs. The earliest explanation common in both arcs is the generalized cause for a character's emotional state. This is process illustrated in figure 3-5.

Between *Rudolph the Red-Nosed Reindeer* and *The Math Aficionado*, the backtracking from the `sad` emotional state stops when it reaches `X is different` in both emotional arcs. This is the earliest, common explanation, with story-specific details abstracted away and, thus, generalizable. Similarly, the backtracking from the `happy` emotional state stops when it reaches `X is different, and X is valuable`. The emotional arcs from both stories are shown in table 3.1. The earliest, common expla-

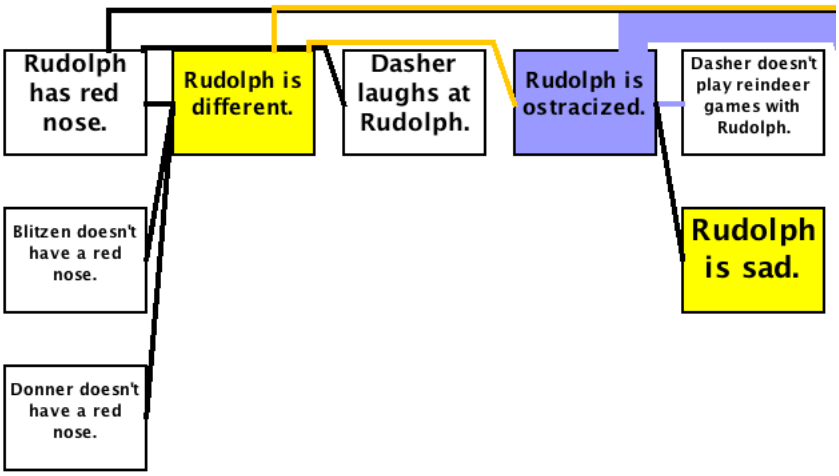

Figure 3-4: The emotional arc for Rudolph's sadness, enlarged from figure 3-2. The endpoint of the emotional arc is an inferred emotional state highlighted in yellow: `Rudolph is sad`. This story element is not explicitly in the story. The beginning of the emotional arc shows the initial cause for the endpoint emotional state: `Rudolph has red nose` while other reindeer don't.

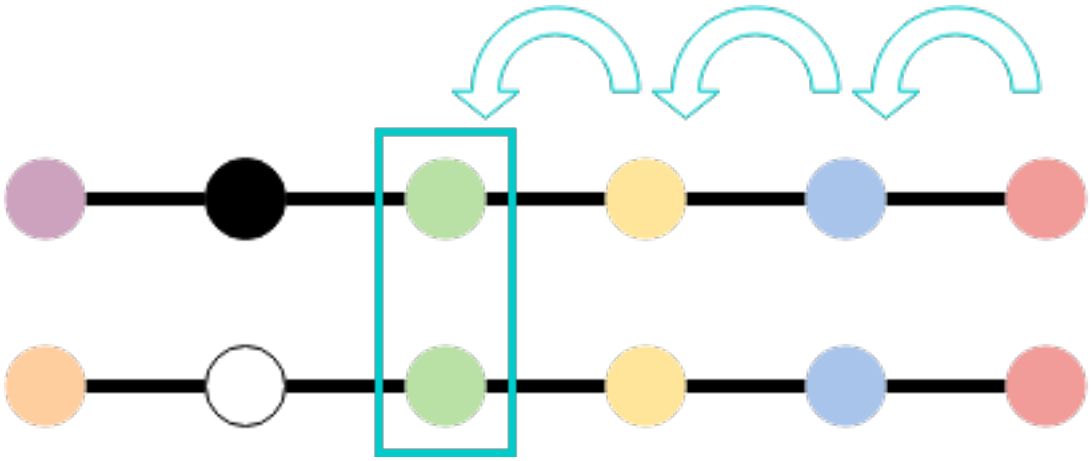

Figure 3-5: MAXIM aligns the emotional arcs in the fable and the past experience to abstract away story-specific details. Starting from the endpoints of the emotional arcs, the system backtracks simultaneously in both arcs until the explanation differs. In this figure, each node is a story element. Matching story elements are indicated using the same color. The emotional arcs diverge before the green node, indicating that the green node contains the general explanation for the emotional state represented using the pink node. The backtracking process stops at the green node.

nations are summarized in table 3.2.

Rudolph the Red-Nosed Reindeer	*The Math Aficionado*
`Rudolph is different because Rudolph has a red nose, Blitzen doesn't have a red nose, and Donner doesn't have a red nose.`	`Tyrion is different because Jamie likes sports, Tyrion likes math, and Tywin likes sports.`
`Rudolph is ostracized, probably because Rudolph is different.`	`Tyrion is ostracized, probably because Tyrion is different.`
`Rudolph is sad because Rudolph is ostracized.`	`Tyrion is sad because Tyrion is ostracized.`
`Rudolph is welcomed, probably because Rudolph is different, and Rudolph is valuable.`	`Tyrion is welcomed, probably because Tyrion is different, and Tyrion is valuable.`
`Rudolph is happy because Rudolph is welcomed.`	`Tyrion is happy because Tyrion is welcomed.`

Table 3.1: Emotional arcs for two emotional states—`X is sad` and `X is happy`—in both the fable and the past experience are shown. For sadness, the earliest explanation common in both stories is being different. For happiness, the earliest explanation common in both stories is the difference being valuable. Note that in the English translation of each inference, the consequence is stated first, followed by the cause. An artifact of the translation process is that the consequence and the cause of an inference are connected by `probably because` if the inference is made from an explanation rule. An explanation rule is applied if there is no other explanation for the the consequence. The conjunction does not indicate probabilistic modeling.

3.3 Grouping Emotional Arcs by Character

Following the *Viewpoint Character Principle*, character arcs are isolated. Emotional states are grouped by their associated characters to isolate character arcs. Each story element is represented by an Entity in Genesis's internal representation, Innerese. The character associated with a story element is determined by examining the subject of the element's Entity.

Rudolph the Red-Nosed Reindeer follows many reindeer, nine to be exact. Parts of Rudolph's and Dasher's character arcs are listed in Innerese in table 3.3:

Having multiple characters in a story implies having emotional arcs belonging to different character, which begs the question of how the process of finding common

Rudolph the Red-Nosed Reindeer	*The Math Aficionado*
`Rudolph is different` → `Rudolph is sad`	`Tyrion is different` → `Tyrion is sad`
`Valuable Rudolph is different` → `Rudolph is happy`	`Valuable Tyrion is different` → `Tyrion is happy`

Table 3.2: Generalized explanations for emotional states in the fable and the past experience. MAXIM concludes being `different` and being `different` but `valuable` are the earliest, common explanations of both character's emotional arcs.

Rudolph	**Dasher**
`(rel have (ent rudolph-5751) (seq roles (fun object (ent nose-5815))))`	`(rel not have (ent dasher-5519) (seq roles (fun object (ent nose-5779))))`
`(rel has-mental-state (ent rudolph-5751) (seq roles (fun object (ent sad-4787))))`	`(rel not play (ent dasher-5519) (seq roles (fun object (ent reindeer_games-6495)) (fun with (ent rudolph-5751))))`

Table 3.3: Parts of Rudolph's and Dasher's character arcs in *Rudolph the Red-Nosed Reindeer*. The character associated with each story element is identified by examining the subject of the Entity.

explanation, described in the previous section, determines which two emotional arcs to align. Emotional arcs should be aligned by character, as well as by the endpoint emotional state. However, in my thesis, alignment is easy because the fable and the past experience are written with the same structure. So, emotional arcs are in order and do not have to be aligned by examining both the associated character and the emotional state.

3.4 Extracting the Lesson Illustrated by a Character

Following the *Reversal of Fortune principle*, MAXIM isolates the central moral challenge by evaluating whether a character's emotional state undergoes an upturn, signalling a motivational story, or a downturn, signalling a cautionary tale. An emotional state is categorized as `POSITIVE`, `NEGATIVE`, or `NEUTRAL`. Table 3.4 shows how a character's emotional state reverses in different types of stories.

The system can be provided with common sense knowledge of the emotional state

	Initial	Final
Motivational Story	NEGATIVE NEUTRAL	POSITIVE
Cautionary Tale	POSITIVE NEUTRAL	NEGATIVE

Table 3.4: Reversal of a character's emotional state in a motivational story and a cautionary tale

category via an input file. To interpret *Rudolph the Red-Nosed Reindeer*, MAXIM is provided with the following knowledge:

Emotional State	Category
happy	POSITIVE
sad	NEGATIVE

From the upturn in Rudolph's emotional state, MAXIM identifies that it is a motivational story.

3.5 Explaining the Moral with Evidence from the Fable

Aside from identifying the moral, the system should also be able to explain it, showing evidence from the fable that supports its conclusion. To this end, MAXIM finds the character's reaction to a challenge. Recall that the system explained the initial emotional state to identify the challenge. The same strategy is employed in this case. The character's reaction to a challenge is embedded in the original cause that resulted in the character's final emotional state.

In Rudolph's case, the explanation for his happy ending is `Valuable Rudolph is different`, as seen in table 3.2. Again, it is important to note that this explanation does not consist of story-specific details. The explanation is abstract and generalizable to other situations.

3.6 Translating from MAXIM's Representation to English Description

Once MAXIM has identified the moral of the story, it translates the moral from its inner representation to an English description using the START parser in Genesis. The central, moral challenge is appended to a template, which differs depending on the type of the character arc, as shown in table 3.5.

	Template
Motivational Story	`It's ok if you <challenge>.`
Cautionary Tale	`It's bad if you <challenge>.`

Table 3.5: The central moral challenge is translated to English and appended to a template, which differs depending on the type of the character arc.

All together, from *Rudolph the Red-Nosed Reindeer*, MAXIM concludes:

`It's ok to be different`

because

`Valuable Rudolph is different`.

3.7 Implications of Emotional Explanation and Alignment

Note that by explaining the change in a character's emotional state, the system can identify the character's challenge without prior knowledge of what constitutes a challenge. The explanation of the initial emotional state shows the challenge that the character faced: Rudolph is sad because other reindeer do not let him join the reindeer games for being the only one with a red nose. On the other hand, the explanation of the final emotional state shows how the character navigated through the challenge: Rudolph is happy because he is invited to the reindeer games after his red nose lit up a foggy Christmas Eve. While the system is provided with common sense knowledge of what it means to be different or valuable, it does not know being different can be a challenge, or the challenge can be overcome by proving the difference can be

valuable. MAXIM arrives at these conclusions by explaining the emotional states and the transitions, in line with the *Emotional Explanation Principle*.

It is also important to reiterate that neither the challenge nor the reaction is an idiosyncratic detail of the story, such as having a red-nose or lighting up a foggy Christmas Eve. By aligning the emotional arcs in a fable and a past experience, MAXIM finds the earliest, common explanation for the emotional states. Thus, the system is able to generalize and find a level of abstraction that facilitates drawing analogies and applying the moral. From Rudolph, MAXIM extracts that being different, which can mean having a red nose when no one else does, can be challenging. However, the difference can be valuable, which can mean being the only reindeer to be able to light up the sky. This is in line with the *Emotional Alignment Principle*

Chapter 4

What Can We Learn from a Reindeer, a Shepherd Boy, a Tortoise, and a Hare?

This chapter shows MAXIM's output for three different fables, along with some discussions about the results. The stories can be found in the appendix. Each moral is followed by an explanation using original story elements.

4.1 *Rudolph the Red-Nosed Reindeer*

Chapter 3: MAXIM explains in detail how the moral of *Rudolph the Red-Nosed Reindeer* is extracted by emotionally explaining and aligning it with *The Math Aficionado*. To reiterate, the system identifies:

```
It's ok if you are different.
Valuable Rudolph is different.
```

4.2 *The Boy Who Cried Wolf*

MAXIM extracts the moral of *The Boy Who Cried Wolf* by emotionally aligning it with *Scheming Ponzi*, a story about a boy named Ponzi who lied about having

stomachaches to get out of school, only to catch stomach flu one day and not be believed by his parents. Both stories involve a character who initially feels entertained that people fell for his lie. Then, the character loses the trust of people around him after lying repeatedly and gets in trouble: one loses his sheep to a wolf and the other is not taken to the hospital to get his stomach flu treated. The system identifies an emotional downturn and finds the common explanation between the two stories, concluding:

```
It's bad if you lie to others.
The Shepherd Boy is in trouble because The villagers ignore Shepherd
    Boy's plea for help, and Trouble is brewing.
```

In the explanation, *trouble is brewing* refers to a wolf appearing, which parallels Ponzi catching a stomach flu. The characters are not quite in trouble yet, if they receive help. However, when their plea for help is ignored, they are truly in trouble.

4.3 *The Tortoise and the Hare*

The Tortoise and the Hare is emotionally aligned with *The Persevering Scientist*, in which one persevering student studies diligently to get better in science while a precocious student becomes complacent and falls behind.

The Tortoise and the Hare is particularly interesting because the story features two contrasting characters who illustrate opposite lessons. With the Tortoise as the viewpoint character, MAXIM identifies an upturn in emotional state and, thus, extracts a motivational moral:

```
It's ok if you are at a disadvantage.
The Tortoise perseveres.
```

On the other hand, with the Hare as the viewpoint character, the system identifies a downturn in emotional state and, thus, extracts a cautionary moral:

```
It's bad if you are at an advantage.
Tortoise is ahead of Hare, probably because Tortoise perseveres.
```
[1]

[1]Note that an artifact of the translation process from Innerese to English is that the consequence

The explanation for the cautionary moral shows a limitation of the system. The explanation revolving around the Tortoise is not convincing. It would be more reasonable to caution against having an advantage, if the explanation shows how this was bad for the Hare. In particular, Hare's complacency cost it the race, and, thus, one should beware not to mismanage an advantage.

This limitation arises because of the way the story is written and the common sense knowledge provided. In particular, the story states that while the Hare naps, the Tortoise gets ahead. People may naturally infer that this also means the Hare is behind, but the system cannot do this unless provided with this common sense knowledge. Similarly, the provided deduction rules are also written from one perspective:

```
If xx is at a disadvantage and xx is ahead of yy, then xx feels
    victorious.
If yy is at an advantage and xx is ahead of yy, then yy feels
    embarrassed.
```

While it is not difficult to provide the knowledge that X being ahead of Y implies Y being behind X, it means the same plot is repeated from different points of view. Instead, it may be beneficial for the system to have the concept of *opposites*. This way, a story does not need to be elaborated to contain repeating story elements that are the opposite sides of the same coin.

and the cause of an inference are connected by `probably because` if the inference is made from an explanation rule. An explanation rule is applied if there is no other explanation for the the consequence. The conjunction does not indicate probabilistic modeling.

Chapter 5

Contributions

5.1 The Next Chapter

I have demonstrated how a generalizable moral can be extracted from a fable. I have begun by interpreting the fable emotionally and aligning it with a past experience consisting of similar emotional arcs. There are many ways to take it to the next level.

Complex Narrative Structure

Currently, MAXIM reads fables in which each character has a single reversal of fortune. However, more complex stories with multiple reversals of fortune may also illustrate morals. For the system to scale up, it should be able identify and evaluate all reversals of fortune within a character arc. More research is needed to determine if a single, overall moral should be extracted, or the morals of each reversal of fortune should influence each other and be adjusted as the story unfolds.

Moral Ambiguity

In this thesis, I categorize a character arc as either a motivational story, with the purpose of providing guidance on how to navigate a difficult situation, or a cautionary tale, with the purpose of warning readers of the consequence of lapse in judgement. Arguably, the world is hardly black and white, and complex stories can illustrate

more ambiguous lessons. To take the MAXIM to the next level, research is needed to determine how to extract types of morals other than `It's ok if you....` and `It's bad if you...`

Applying the Moral

While an English description of the moral is insightful, it would be valuable to see the morals applied. To feed the learned moral back into the system, the moral can be converted into an explanation rule and provided to the system as common sense knowledge. Fables are often read to instill ethical values in children. Research into how lessons from fables impact ethical choices may provide fascinating insights. For instance, if a situation resembles bits and pieces of different fables, how can a person leverage the morals of each fable to navigate the situation?

5.2 Contributions

In this thesis, I have taken a step toward understanding an aspect of human intelligence.

1. **I proposed four principles for extracting a generalizable moral from a fable.**

   Viewpoint Character Principle: the moral of the story depends on the point of view.

   Reversal of Fortune Principle: to find the moral of the story, look for a reversal of fortune. The change of a character's fortune for the better or worse signals the consequence of the character's choice when faced with a challenge.

   Emotional Explanation Principle: to find a reversal of fortune, look for emotional arcs by explaining the emotional states and their transitions in the story.

Emotional Alignment Principle: to find the right level of abstraction for the moral, align the fable and a past experience with similar emotional arcs.

These principles are formed by finding some commonalities among fables and empirically grounded in psychology research. Together, they inform how to identify the challenge specific to a story and to generalize it to be applicable to other situations. Though internalized by adults so that they feel intuitive, they are learned, perhaps unconsciously, by children.

2. **I evaluated and demonstrated the fundamental role of emotion in learning a moral.** By emotionally explaining and aligning a story with a past experience to learn a generalizable moral, I illustrated that emotional empathy is the heart of a story, not an afterthought.

3. **I built MAXIM**, a new module in Genesis. From a fable and a description of a past experience written in English, MAXIM extracts and generalizes the moral of the fable and explains it in English. MAXIM takes Genesis to the next level by demonstrating that an artificial intelligence system can learn from a story like a human. Stories are often told to instill moral values in children. MAXIM paves the way for modeling moral development empowered by stories.

4. **I demonstrated that an artificial intelligence system can identify the moral of a story, similar to how a human does.** MAXIM follows a humanly viable computational process. I have demonstrated the system at work with three different fables, each of which has a different narrative structure. These fables illustrate both motivational and cautionary morals.

Appendix A

The Genesis Enterprise

MAXIM is implemented on top of the Genesis Story Understanding System. To understand the inner workings of MAXIM, it is necessary to first understand the inner workings of Genesis. This section describes the Genesis Enterprise, an endeavor founded the view that story understanding, made possible by the uniquely human, merge-enabled inner language, provides the substrate for other aspects of human intelligence, such as problem solving. The Genesis Story Understanding System has been developed to build a computational model of the uniquely human story understanding capability.

A.1 Story Understanding is Uniquely Human

Over the course of evolution, how is it that only we—humans—developed the ability to understand stories? The followings sections review arguments for what sets us apart from other species.

A.1.1 Merge Makes Us, and Only Us, a Symbolic Species

Berwick and Chomsky argue that *Merge* is what sets us apart from other primates. Biologically, this operation is enabled by the closing of an anatomical loop in the human brain that is nearly complete in the brains of other primates.

***Merge*:** an operation that takes two objects already constructed, call them X and Y, and forms from them a new object that consists of the two unchanged, hence simply the set with X and Y as members. Provided with conceptual atoms of the lexicon, the operation Merge, iterated without bound, yields an infinity of digital, hierarchically structured expressions [1].

The *Merge* operation give us, and only us, an inner language to represent the world around us. Our internal representation enables us to describe what we see and interpret ideas. Using our inner language, we solve problems, create new ideas, and understand stories. When we communicate using words, we externalize by translating our inner language to a natural language and internalize by doing the reverse. Having a merge-enabled inner language is what it means to be symbolic.

A.1.2 Being Symbolic Enables Us to Understand Stories

We use our inner language to form inner stories, which can be internalized and externalized through different mediums, such as speech, text, and sign.

An inner story: A collection of complex, highly nested symbolic descriptions of properties, relations, actions, and events, usefully connected with constraints such as causal, enablement, and time constraints [14]

Unlike stories exchanged between people using natural language, inner stories do not need to adhere to some narrative structure or follow a protagonist. A set of instruction for changing a cell phone battery, or a recipe, can be an inner story.

A.1.3 Story Understanding is the Substrate of Human Intelligence

Our ability to understand stories rests on the uniquely human, merge-enabled inner language. Stories play integral roles across cultures and societies. We tell stories to entertain; we tell stories to memorialize events; we tell stories to teach life lessons.

Looking beyond the prevalence and significance of stories, it is worth remembering that only we tell stories.

> **The Strong Story Hypothesis:** The mechanisms that enable humans to tell, to understand, and to recombine stories separate our intelligence form that of other primates [13]

So, if we are to develop a computational model of human intelligence, then we must model how we understand stories.

A.2 The Genesis Story Understanding System

The Genesis Enterprise is undertaken with the principle that we cannot thoroughly learn the details of computation unless we implement a system that models the behavior that we want to understand. The Genesis Story Understanding System models various aspects of how people understand stories, such as aligning different stories, interpreting stories with cultural biases, and drawing analogies to similar stories. Genesis is developed with the computational-imperative principle in mind.

> **The Computational-Imperative Principle:** Any model of human intelligence should introduce only computational capabilities that enable observed behaviors without enabling unobserved behaviors [14]

Crucially, Genesis aims to model what humans do naturally, and nothing beyond, lest we model something so general that it explain more than computations enabled by human intelligence.

MAXIM is added as a module in Genesis. The following chapters describe the implementation that MAXIM is built upon.

A.2.1 START: English to Innerese

Genesis's Story Processor reads stories written in English and uses the START parser (Katz, 1997) to translate to Innerese, its inner language. Innerese consists of entities,

functions, relations, and sequences. From the perspective of Java language, in which Genesis is implemented, Entity is a parent class. Function inherits Entity, Relation inherits Function, and Sequence inherits Entity.

An **entity** is an atomic unit of the Innerese. Each entity consists of a unique name and a bundle of threads, each of which provides a meaning of the entity derived from WordNet [5]. An example of an entity is a single object, such as a *mug*.

A **function** is an entity with a subject, which is also an entity. Functions represent Jackendoff's places and path elements [7], such as the *inside of a mug*, in which the mug is the subject.

A **relation** is a function with an object, which is also an entity. It shows how the subject relates to the object, such as *coffee inside the mug*.

A **sequence** is an ordered set of entities, such as *Ben Bitdiddle drank all the coffee inside the mug*.

A.2.2 Inferences Beyond Explicit Story Elements

We need to read between the lines, considering that stories often imply meanings and consequences rather than directly expressing them. As we read, we use our common sense knowledge to draw inferences between elements in the story. Sometimes, part of an inference, such the consequence of some event, is not present in the story. Like people, Genesis's Story Processor can make inferences beyond explicit story elements based on its common sense knowledge. When inferences are made, new story elements are injected into the story if not already present. Otherwise, disjoint, existing story elements are connected. When new story elements are injected, the Story Processor searches through its collection of common sense knowledge to check if further inferences can be made. Genesis's common sense knowledge are provided as different types of rules rules, some of which used to understand the fables in this thesis are detailed below.

Deduction Rule

```
If X is ostracized, then X is sad.
```

A deduction rule infers the consequent that directly follows the antecedent. From the above rule, if X is ostracized, then it follows that X is sad.

Explanation Rule

```
If X is different, then X may be ostracized.
```

An explanation rule is applied when an event does not have an explanation otherwise. In other words, the antecedent does not always lead to the consequent. In the example above, if there is no other explanation for why X is ostracized, and the story states that X is different, then a causal connection is made between these two story elements.

Abduction Rule

```
If X doesn't play reindeer games with Y, then Y must be ostracized
```

An abduction rule explains the cause for an event, even if the cause is not already in the story or if there are alternative causes. Unlike the other rules, a cause is the consequent in an antecedent rule, and a consequence is the antecedent. The above rule states Y being ostracized may be one of the causes that X doesn't play reindeer games with Y.

Appendix B

Empirical Grounding for the Guiding Principles

The guiding principles for extracting a generalizable moral of a story are empirically grounded in psychology research. This section describes why it is important to explain and align stories, and why, specifically, it is important to do so emotionally.

B.1 The Case for Explanation and Alignment with Past Experience

B.1.1 Self-Explaining and Aligning the Story Facilitates Abstraction

Research has shown that prompting children to explain key events of a fable facilitates abstracting the moral of the story [11]. In Walker and Lombrozo's experiment, children between 5- and 6-years-old read fables and were asked to explain why certain key events occurred. The results showed that children who were prompted to explain were less likely to focus on idiosyncrasies of the story and more likely to extract a generalizable lesson.

One of the fables used in the experiments is about a short person who is first

ostracized then accepted in a town full of tall people, children who were prompted to explain why the short person was sad and then happy at two points in the story were more likely to conclude that it's ok to be different. On the other hand, children in the control group were more likely to focus on the story-specific detail: tall people will play with a short person.

The experiment also showed that it was not enough to simply inform children of the reasons why key events occurred; instead, it was crucial for children to generate the explanation by themselves. This suggests that it is not sufficient to simply draw children's attention to key events associated with the moral.

The authors hypothesize that in trying to explain why an event occurred, children call upon their prior knowledge to find an explanation that is broad [12] and simple [2] to account for as much of the story as possible without introducing additional constraints. It is possible that children align the fable at hand with a past experience.

In the above example, in explaining why a short person was sad, children may recall a time when they encountered a person who was also ostracized because of his difference, which may be some characteristic other than being short. Aligning the fable and the past experience reveals that a broad and simple explanation is difference, not being short in particular, can lead to ostracism.

While the research sheds valuable insight on the computational steps in abstracting away story-specific details to extract a generalizable moral, it is less clear what constitutes a key event that should be explained. This will be discussed in the next section B.2: The Case for Emotion.

B.1.2 Self-Explaining Facilitates Learning in General

Walker and Lombrozo's experiment and hypothesis are consistent with the broader line of research that proposes self-explaining facilitates learning in general. Chi argues that the process of self-explanation repairs our mental models when there is a conflict between our prior knowledge and new, incoming information [3]. By explaining the new, incoming information to ourselves, we generate new inferences with which we use to fill in the gaps in our existing mental model. This extends beyond the previous line

of thought that self-explanation is simply a process of inference generation, during which we elaborate on the new information using our prior knowledge [4]. In particular, the number of sound inferences generated in the process of self-explanation is not dependent on the amount of prior knowledge. Greater amount of prior knowledge does not always yield more generated inferences. Instead, we generate inferences to revise our prior knowledge with new information.

B.2 The Case for Emotion

B.2.1 Readers Naturally Infer a Character's Emotional States

Research has shown that readers naturally infer a fictional character's emotional state [6]. They do this without being prompted, or even when the story contains no mention of emotions. In a series of experiments, Gernsbacher et al. presented subjects with stories that contained concrete actions but no emotions. After reading each story, subjects read a target sentence that contained a word that either matched or mismatched the implied emotional state of a character.

For instance, subjects read a story about a protagonist who stole from a store during his friend's shift, leading to his friend getting fired. The results showed that subjects read target sentences more slowly if they contained mismatching emotional state—pride—than if they contained matching emotional state—guilt. This suggests that in the process of reading, a reader builds a mental representation of the story that includes inferences of the character's emotional state, and the reader does this naturally without being prompted.

B.2.2 Readers Feel a Character's Pain

Our mental representation of stories automatically include the characters' emotional states. This raises the question of how we leverage this knowledge. What role does our emotional knowledge play in aligning two stories?

Thagard argues that analogies can be made from literature to real life stories if

emotional structures align [10]. When reading a work of fiction, readers empathize with a character's journey, sharing the character's pain or joy as readers imagine themselves undergoing a similar adventure. He argues that emotion is a crucial component of an allegory, a piece of work typically concealing a hidden message under salient story-specific details. As an example, he aruges that George Orwell's *Animal Farm* is a successful allegorical critique on the rise of communism in the Soviet Union because the story invokes similar emotions in readers as that period in history.

Thagard's argument complements Lehnert's proposition that narratives are organized around a reader's knowledge about the character's emotional state, rather than the character's goals or strategies [8]. An event can cause a **positive emotional reaction** in a character if the event pleases the character, **negative** if displeases, and **neutral** if no effect. Two emotional states comprise a primitive plot unit, multiple of which can be combined to form a complex plot unit, which is, in other words, a narrative.

All together, Thagard's and Lehnert's works argue that a story is written to revolve around emotional states of characters, which readers leverage to draw analogies to other stories. These ideas inform that the key events that should be explained are character's emotional states, and a fable and a past experience should be aligned by emotional states.

Appendix C

Stories

C.1 "It's ok if you are different"

C.1.1 *Rudolph the Red-Nosed Reindeer*

```
Santa is a person.


Dasher is a reindeer.
Dancer is a reindeer.
Prancer is a reindeer.
Vixen is a reindeer.
Comet is a reindeer.
Cupid is a reindeer.
Donner is a reindeer.
Blitzen is a reindeer.
Rudolph is a reindeer.


Dasher doesn't have a red nose.
Dancer doesn't have a red nose.
Prancer doesn't have a red nose.
Vixen doesn't have a red nose.
Comet doesn't have a red nose.
```

Cupid doesn't have a red nose.

Donner doesn't have a red nose.

Blitzen doesn't have a red nose.

Rudolph has a red nose.

Dasher laughs at Rudolph because Rudolph has a red nose.

Dancer laughs at Rudolph's red nose.

Prancer laughs at Rudolph's red nose.

Vixen laughs at Rudolph's red nose.

Comet laughs at Rudolph's red nose.

Cupid laughs at Rudolph's red nose.

Donner laughs at Rudolph's red nose.

Blitzen laughs at Rudolph's red nose.

Dasher doesn't play reindeer games with Rudolph.

Dancer doesn't play reindeer games with Rudolph.

Prancer doesn't play reindeer games with Rudolph.

Vixen doesn't play reindeer games with Rudolph.

Comet doesn't play reindeer games with Rudolph.

Cupid doesn't play reindeer games with Rudolph.

Donner doesn't play reindeer games with Rudolph.

Blitzen doesn't play reindeer games with Rudolph.

Rudolph is sad.

Santa can't guide the sleigh because he can't see clearly on a foggy Christmas Eve.

Rudolph's red nose can light up the sky.

Santa asks Rudolph to guide the sleigh because Rudolph has a red nose.

Dasher apologizes to Rudolph for laughing at his red nose.

Dasher plays reindeer games with Rudolph.

Rudolph is happy.

C.1.2 *The Math Aficionado*

```
Tywin is a person.
Tyrion is a person.
Jamie is a person.
Cersei is a person.

Tyrion likes math.
Tywin, Jamie, and Cersei like sports.

Tywin and Cersei laugh at Tyrion because Tyrion likes math.
Tywin, and Cersei don't sit with Tyrion at lunch.
Tyrion is sad.

Jamie can't integrate well.
Tyrion can integrate well.
Jamie asks Tyrion to teach him integration because Tyrion likes math.

Jamie sits with Tyrion at lunch.
Tyrion is happy.
```

C.1.3 Common Sense Knowledge

```
xx, yy, and zz are persons.

// Prediction rules
If xx has red nose and yy doesn't have red nose and zz doesn't have
    red nose, then xx is different.
If xx likes math and yy likes sports and zz likes sports, then xx is
    different.
If xx can't guide the sleigh and xx asks yy to guide the sleigh, then
    yy is valuable.
If xx can't integrate well and xx asks yy to teach him integration,
```

```
    then yy is valuable.
If xx is ostracized, then xx is sad.
If yy is welcomed, then yy is happy.

// Explanation rules
If xx is different, then xx may be ostracized.
If xx is different and xx is valuable, then xx may be welcomed.

// Abduction rules
If xx doesn't play reindeer games with yy, yy must be ostracized.
If xx doesn't sit with yy at lunch, yy must be ostracized.
If xx plays reindeer games with yy, yy must be welcomed.
If xx sits with yy at lunch, yy must be welcomed.
```

C.2 "It's bad if you lie"

C.2.1 *The Boy Who Cried Wolf*

```
Shepherd Boy is person.
Villager is a person.

The Shepherd Boy tends sheep near the village.
The Shepherd Boy is bored because he doesn't have much to do.

One day, the Shepherd Boy has an idea.

The Shepherd Boy yells wolf.
The Villagers come to the Shepherd Boy's rescue.
The Villagers only see sheep.
The Shepherd Boy feels entertained because he lied to the Villagers
    and the Villagers believed him.
The Villagers are angry.
```

```
Then, a few days pass.

The Shepherd Boy yells wolf again.
The Villagers come to the Shepherd Boy's rescue again.
The Villagers only see sheep again.
The Shepherd Boy feels more entertained because he repeatedly lied to
    the Villagers and the Villagers believed him.
The Villagers are angrier.

Then, a more few days pass.

A wolf appears.
The Shepherd Boy shouted "wolf".
The Villagers ignore the Shepherd Boy's plea for help.
The wolf eats the Shepherd Boy's sheep.
The Shepherd Boy feels frightened.
```

C.2.2 *Scheming Ponzi*

```
Ponzi is person.
Parent is a person.
Doctor is a person.

Ponzi does not like school.

One day, Ponzi has an idea.

Ponzi complains that his stomach hurts to his parent.
Ponzi's parent take him to the hospital.
The doctor runs some tests.
Ponzi is healthy.
Ponzi feels entertained because Ponzi lied to his parent and his
```

```
   parent believed him.
Ponzi's parent is angry.

Then, a few weeks pass.

Ponzi complains again that his stomach hurts to his parent.
Ponzi's parent take him to the hospital again.
The doctor runs some tests again.
Ponzi is still healthy.
Ponzi feels more entertained because Ponzi repeatedly lied to his
   parent and his parent believed him.
Ponzi's parent is angrier.

Then, a more few weeks pass.

Ponzi catches the stomach flu.
Ponzi says his stomach is in pain.
Ponzi's parent ignores Ponzi's plea for help.
Ponzi goes to school.
Ponzi throws up and has a fever.
Ponzi feels frightened.
```

C.2.3 Common Sense Knowledge

```
xx and yy are persons.

// Prediction rules
If xx yells wolf and yy only sees sheep, then xx lies to yy.
If xx is healthy and xx complains that his stomach hurts to yy, then
   xx lies to yy.
If xx yells wolf again and yy only sees sheep again, then xx
   repeatedly lies to yy.
If xx is still healthy and xx complains again that his stomach hurts
```

```
   to yy, then xx repeatedly lies to yy.
If a wolf appears, then trouble is brewing.
If xx catches the stomach flu, then trouble is brewing.
If trouble brews and xx ignores yy's plea for help, then yy is in
   trouble.
If yy lies to xx and xx believes yy, then xx feels angry.
If yy repeatedly lies to xx and xx believes yy, then xx feels angrier.
If xx is in trouble, then xx feels frightened.

// Explanation rules
If xx lies to yy, then yy may believe xx.
If xx repeatedly lies to yy, then yy may distrust xx.

// Abduction rules
If xx comes to yy's rescue, xx must believe yy.
If xx takes yy to the hospital, xx must believe yy.
If xx ignores yy's plea for help, xx must distrust yy.
```

C.3 "It's ok if you are at a disadvantage" and "It's bad if you are at an advantage"

C.3.1 *The Tortoise and the Hare*

```
Tony the Tortoise and Hare are persons.

Hare can run quickly.
Tony the Tortoise runs slowly.

Hare mocks Tony the Tortoise's slow pace.
Hare says "how do you get anywhere".
Tony the Tortoise replies "I might just beat you in a running race".
Hare is amused at the Tony the Tortoise's challenge because Hare can
```

```
    run quickly.

Hare and Tony the Tortoise run a race.
Tony the Tortoise starts crawling slowly.
Hare quickly sprints past Tony the Tortoise.
Hare is ahead of Tony the Tortoise.

Hare takes a nap in the middle of the race.
Tony the Tortoise doesn't stop crawling.

Then, an hour passes.
Tony the Tortoise crawls past Hare.
Tony the Tortoise is ahead of Hare.

Then, another hour passes.
Hare wakes up and sees Tony the Tortoise near the finish line.
Hare sprints to the finish line but cannot catch up.
Tony the Tortoise wins the race.
```

C.3.2 *The Persevering Scientist*

```
Rick and Morty are persons.

Rick is good in science.
Morty is terrible in science.

Rick mocks Morty's lack of scientific knowledge.
Rick says "how can you not understand science".
Tortoise replies "I might just be a better scientist than you some day
    ".
Rick is amused at the Morty's challenge because Rick is good at
    science.
```

```
Rick and Morty take a science exam.
Rick scores above Morty.
Rick is ahead of Morty.

Rick plays all day long.
Morty doesn't stop studying.

Then, a few months pass.
Rick and Morty take another science exam.
Morty scores above Rick.
Morty is ahead of Rick.
```

C.3.3 Common Sense Knowledge

```
xx and yy are persons.

// Prediction rules
If xx can run quickly, then xx is at an advantage.
If xx is good in science, then xx is at an advantage.
If xx runs slowly, then xx is at a disadvantage.
If xx is terrible in science, then xx is at a disadvantage.
If xx is at an advantage and xx is ahead of yy, then xx feels smug.
If yy is at a disadvantage and xx is ahead of yy, then yy feels
    disheartened.
If xx is at a disadvantage and xx is ahead of yy, then xx feels
    victorious.
If yy is at an advantage and xx is ahead of yy, then yy feels
    embarrassed.

// Explanation rules
If xx is at an advantage and yy is at a disadvantage, then xx may be
    ahead of yy.
If xx perseveres, then xx may be ahead of yy.
```

```
// Abduction rules
If xx take a nap in the middle of the race, xx must be complacent.
If xx plays all day long, xx must be complacent.
If xx doesn't stop crawling, xx must persevere.
If xx doesn't stop studying, xx must persevere.
```

Bibliography

[1] Robert C Berwick and Noam Chomsky. *Why only us: Language and evolution*. MIT press, 2016.

[2] Elizabeth Baraff Bonawitz and Tania Lombrozo. Occam's rattle: Children's use of simplicity and probability to constrain inference. *Developmental psychology*, 48(4):1156, 2012.

[3] Michelene TH Chi. Self-explaining expository texts: The dual processes of generating inferences and repairing mental models. *Advances in instructional psychology*, 5:161–238, 2000.

[4] Michelene TH Chi, Nicholas De Leeuw, Mei-Hung Chiu, and Christian LaVancher. Eliciting self-explanations improves understanding. *Cognitive science*, 18(3):439–477, 1994.

[5] Christiane Fellbaum. Wordnet: An electronic lexical database and some of its applications, 1998.

[6] Morton Ann Gernsbacher, H Hill Goldsmith, and Rachel RW Robertson. Do readers mentally represent characters' emotional states? *Cognition & Emotion*, 6(2):89–111, 1992.

[7] Ray Jackendoff. *Semantics and cognition*, volume 8. MIT press, 1983.

[8] Wendy G Lehnert and Elaine W Vine. The role of affect in narrative structure. *Cognition and emotion*, 1(3):299–322, 1987.

[9] Darcia Narvaez, Tracy Gleason, Christyan Mitchell, and Jennifer Bentley. Moral theme comprehension in children. *Journal of Educational Psychology*, 91(3):477, 1999.

[10] Paul Thagard. The brain is wider than the sky: Analogy, emotion, and allegory. *Metaphor and Symbol*, 26(2):131–142, 2011.

[11] Caren M Walker and Tania Lombrozo. Explaining the moral of the story. *Cognition*, 167:266–281, 2017.

[12] Caren M. Walker, Tania Lombrozo, Joseph J. Williams, Anna N. Rafferty, and Alison Gopnik. Explaining constrains causal learning in childhood. *Child Development*, 88(1):229–246.

[13] Patrick Henry Winston. The strong story hypothesis and the directed perception hypothesis. In *2011 AAAI Fall Symposium Series*, 2011.

[14] Patrick Henry Winston and Dylan Holmes. The genesis enterprise: Taking artificial intelligence to another level via a computational account of human story understanding. Technical report, 2018.