COMPUTATIONAL MODELS OF HUMAN INTELLIGENCE

CMHI Report Number 1 **15 December 2018**

The Genesis Enterprise: Taking Artificial Intelligence to another Level via a Computational Account of Human Story Understanding

Patrick Henry Winston and Dylan Holmes

Computer Science and Artificial Intelligence Laboratory
Center for Brains, Minds, and Machines

We propose to develop computational accounts of human intelligence and to take intelligent systems to another level using those computational accounts.

To develop computational accounts of human intelligence, we believe we must develop biologically plausible models of human story understanding, and then use those models to implement story-understanding systems that embody computational imperatives.

We illustrate our approach by describing the development of the Genesis Story Understanding System and by explaining how Genesis goes about understanding short, up to 100-sentence stories, expressed in English. The stories include, for example, summaries of plays, fairy tales, international conflicts, and Native American creation myths.

Genesis answers questions, interprets with controllable allegiances and cultural biases, notes personality traits, anticipates trouble, measures conceptual similarity, aligns stories, reasons analogically, summarizes, tells persuasively, composes new stories, and performs story-grounded hypothetical reasoning.

We explain how we ensure that work on Genesis is scientifically grounded; we identify representative questions to be answered by our Brain and Cognitive Science colleagues; and we note why story understanding has much to offer not only to Artificial Intelligence but also to fields such as business, economics, education, humanities, law, neuroscience, medicine, and politics.

Keywords: computational models of human intelligence; story understanding; computational imperatives; inference reflexes; concept recognition; Genesis story-understanding system.

The Genesis Enterprise: Taking Artificial Intelligence to another Level via a Computational Account of Human Story Understanding

Patrick Henry Winston and Dylan Holmes

Computer Science and Artificial Intelligence Laboratory
Center for Brains, Minds, and Machines

15 December 2018

Abstract

We propose to develop computational accounts of human intelligence and to take intelligent systems to another level using those computational accounts.

To develop computational accounts of human intelligence, we believe we must develop biologically plausible models of human story understanding, and then use those models to implement story-understanding systems that embody computational imperatives.

We illustrate our approach by describing the development of the Genesis Story Understanding System and by explaining how Genesis goes about understanding short, up to 100-sentence stories, expressed in English. The stories include, for example, summaries of plays, fairy tales, international conflicts, and Native American creation myths.

Genesis answers questions, interprets with controllable allegiances and cultural biases, notes personality traits, anticipates trouble, measures conceptual similarity, aligns stories, reasons analogically, summarizes, tells persuasively, composes new stories, and performs story-grounded hypothetical reasoning.

We explain how we ensure that work on Genesis is scientifically grounded; we identify representative questions to be answered by our Brain and Cognitive Science colleagues; and we note why story understanding has much to offer not only to Artificial Intelligence but also to fields such as business, economics, education, humanities, law, neuroscience, medicine, and politics.

Keywords: computational models of human intelligence; story understanding; computational imperatives; inference reflexes; concept recognition; Genesis story-understanding system

1 Vision

Most of today's AI research focuses on statistical mechanisms associated with the field of Machine Learning. Those statistical mechanisms exhibit various kinds of perceptual intelligence, often impressively, but shed little light on aspects of intelligence that are uniquely human. We believe that tomorrow's AI will focus on an understanding of our uniquely human intelligence emerging from discoveries on par with the discoveries of Copernicus about our universe, Darwin about our evolution, and Watson and Crick about our biology. Those cognitive mechanisms will take to another level applications aimed at reasoning, planning, control, and cooperation.

Tomorrow's AI applications will astonish the world because they will think and explain themselves, just as we humans think and explain.

To develop an account of our uniquely human intelligence, we believe we must answer two key questions: first, what computational competences are uniquely human; and second, how do the uniquely human competences support and benefit from the computational competences we share with other animals.

Our answer to the uniquely human question is that we became the symbolic species and that becoming symbolic made it possible to become the story-understanding species. Our answer to the support-and-benefit question is that our symbolic competence, and the story-understanding competence that it enables, could not have evolved without myriad elements already in place.

Our position on the uniquely-human question—that we are symbolic story understanders—takes us out of the mainstream, even among those who study story understanding. As explained in section 7.2, we believe what should be studied is not how problem solving enables story understanding, but rather how story understanding provides the substrate for problem solving and everything enabled by problem solving. As explained in section 6.2, we believe that no neural net can have humanlike intelligence without the equivalent of a symbolic abstraction layer.

Because our position is controversial, we devote the rest of this section to explaining in more detail how we decided to focus our research on story understanding. We start by explaining what we mean when we say we are symbolic, how being symbolic enables story understanding, how we approach our study of story understanding, and how developing an account of human story understanding causes us to focus on behaviors of interest, rather than on statistical measures appropriate to application evaluation. Our work on the support-and-benefit question constitutes another story that is not yet ready to be told.

1.1 Merge is the *sine qua non* of being symbolic

We begin by noting recent work by Robert Berwick and Noam Chomsky. In *Why Only Us* (2016), they argue that only humans have what they call *merge*, an operation that combines two expressions to make a larger expression without disturbing the two merged expressions. Berwick and Chomsky emphasize that having merge is an incremental evolutionary step, a step that requires only the completion of an anatomical loop that is almost complete in other primates.

While completion of an anatomical loop is not a big change anatomically, we hypothesize that it enables a giant increase in intelligence, because merge gives us, and only us, an *inner language* with which we build complex, highly nested symbolic descriptions of classes, properties, relations, actions, and events. When we write that we are symbolic, we mean that we have a merge-enabled inner language.

Using our inner language, we can record, for example, that a hawk is a kind of bird, that hawks are fast, that a particular hawk is above a field, that the hawk is hunting, that a squirrel appears, and that John thinks the hawk will try to catch the squirrel.

Other animals seem to have internal representations of some aspects of the outside world, but they seem incapable of constructing complex, highly nested symbolic descriptions. Recent reexamination of work with chimpanzees, for example, shows that chimpanzees do not have humanlike compositional abilities. Charles Yang, in a seminal study of child and chimpanzee corpora, has noted that young children generate novel combinations of words very freely, but Nim Chimpsky, the famous chimpanzee who was exposed to American Sign Language, never provided evidence, via signing, that suggested he had a merge-enabled compositional capability (2013). Evidently, chimpanzees have some ability to understand the names of things and memorize sign sequences, but they do not express via their signing any indication that they have a merge-enabled inner language of complex, highly nested symbolic descriptions.

1.2 Being symbolic is the *sine qua non* of story understanding

We believe that our merge-enabled inner language, which seems to have emerged only about 80,000 years ago (Tattersall, 1998, 2010, 2012), made possible another distinguishing competence: we connect the complex, highly nested symbolic descriptions with various sorts of constraints, including, for example, causal, means-ends, enablement, and time constraints. With such constraints, we form even more complex and highly nested symbolic descriptions.

> **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.

Note that we exclude what others would include. We have no doubt that rats and other animals remember useful sequences, and we have no objection to calling those sequences inner stories. However, when we refer to an *inner story*, we refer to a story expressed in an inner, merge-enabled language that rats, dogs, and chimpanzees either do not have or do not have on anything close to human level. Note also that we include what some narratologists would exclude. There is no requirement for an inner story to have, for example, a narrative arc; or a beginning, middle, and end; or even what Livia Polanyi would call a point (1989).

1.3 Story understanding is the *sine qua non* of human intelligence

Somehow, we developed the means to externalize our inner stories into outer communication languages and to internalize stories presented to us in those outer communication languages. Being social animals, we started telling each other stories.

Understanding the stories we tell each other has immense consequences. Education includes listening to fairy tales, many of which aim to scare us into behaving properly. Later, we acquire precedent stories packaged up in subjects such as history, literature, law, medicine, business, engineering, science, and religion. We learn how to think by acquiring the skills involved in deploying precedent stories. We learn how to find and combine fragments drawn from precedents to compose new stories creatively.

We believe such story-understanding competences are unique to our species because they ultimately depend on having a merge-enabled inner language. Believing such story-understanding competences to be unique to our species led Winston to propose the following hypothesis:

> **The Strong Story Hypothesis:** The mechanisms that enable humans to tell, to understand, and to recombine stories separate our intelligence from that of other primates (Winston, 2011).

Because so much seems to rest on merge-enabled, uniquely human story understanding, we believe that if there is to be a computational account of human intelligence, we must catalog, characterize, and explain our basic inner-story competences. Then, because we cannot fully understand the computations involved until we implement systems that embody biologically plausible models of the inner-story competences.

Note that we have no desire to be human chauvinists. We make no claim that human intelligence is the only kind of intelligence; our only claim is that our inner-story competences give us a unique kind of highly enabling intelligence. We acknowledge that animals from crows to chimpanzees exhibit impressive capabilities, including the use of tools, and we agree that those capabilities are evidence of various kinds of intelligence.

we subscribe to Marvin Minsky's notion of *suitcase* words; that is, words such as *intelligence*, *creativity*, and *consciousness*, are labels attached to so many different meanings they are like giant suitcases, so big you can stuff just about anything into them (1988; 2006).

1.4 We start by specifying the behavior we want to understand

As we start work on modeling an aspect of story understanding, we first define precisely the story-understanding behavior we are trying to understand.

In building the Genesis model, we study the computations required to translate stories of up to 100 sentences, expressed in simple English, into inner stories. Then, we study the computations required to use the inner stories to, for example, answer questions, describe conceptual content, summarize, compare and contrast, react with cultural biases, instruct, reason hypothetically, solve problems, and find useful precedents (Winston, 2011, 2012a,b).

An *outer story* is anything that produces an inner story. Our few to 100-line stories, expressed in English, are outer stories. We have worked with outer stories drawn from various genres, including play plots, conflict accounts, creation myths, news reports, repair recipes, and how-to-repair-it instructions.

We focus our attention on outer stories expressed in English, so when we use the word *story*, we refer to outer stories expressed in natural language. Similarly, when we refer to an inner story, we mean an inner story produced from natural-language input.

We note, however, that an outer story may consist entirely of visual material such as found in pictures, drawings, diagrams, graphs, cartoons, and performance.

Evelina Fedorenko and Rosemary Varley note that subjects who have lost their ability to process natural language still play chess and show other signs of thinking (2016). We conjecture that a relativehly intact inner-language story-understanding competence still functions even though the cortical mechanisms that externalize to and internalize from English and other natural languages are gone.

1.5 We formulate computational problems, posit representations, and build

Following David Marr (1982), when we identify a particular story-understanding behavior we want to understand, we formulate computational problems and posit representations that expose the constraints needed to solve those problems.

Then, we build. We build because we believe that we have not understood a competence on a deep level until we can develop and implement models that manifest the understood competence. By building our Genesis story-understanding system, we ensure that we develop models that are precise, testable, and composable. Building helps us uncover questions otherwise easily missed. Successful building marks the genesis of understanding.

1.6 We adhere to computational imperatives

Because we aim to develop a computational account of human intelligence, we introduce no representation, no constraint, no method, no architectural element, without a *computational imperative* associated with a human behavior. That is, nothing goes into Genesis unless Genesis *needs* it (Marr, 1977). And of course, nothing goes in unless it seems biologically plausible.

The computational-imperative principle: any model of human intelligence should introduce only computational capabilities that enable observed behaviors without enabling unobserved behaviors.

One example of a computational imperative, from our earliest work on Shakespearean plot summaries, is the use of *explanation rules* to account for our human tendency to look for cause: we understand that Macduff may kill Macbeth because Macbeth angered Macduff, even though the causal link is not mentioned, and even though Macduff does not always kill those who anger him.

A second example is the introduction of the *unknowable leads-to relation* to account for the fact that we can acknowledge causal links even in the absence of detailed understanding. We first came across this need when working with Native American Crow creation myths (Yarlott, 2014), which often express explicit unknowables explicitly: "Old Man Coyote made the world from a handful of mud and you will never understand how." Then, once noted, we now find unknowable leads-to relations in all sorts of stories.

A third example is the introduction of *culturally-specific mental models* to account for anthropological variations in story understanding: people from China tend to explain violence in terms of situations that lead to violent behavior, whereas people from American cultures tend to explain violence in terms of character traits that lead to violent behavior (Morris and Peng, 1994; Awad, 2013).

1.7 The computational-imperative principle promotes science and supports engineering

Because we are primarily motivated by our passion for developing a computational account of human intelligence, we naturally aspire to be sure our work is scientifically grounded. When we ask ourselves the is-it-science question, a question often asked by critics of work in Artificial Intelligence, we first think about falsifiability and then consider other ways in which scientific accounts are evaluated.

1.8 What about falsifiability?

The behavior we are trying to explain is story understanding, and that requires hypotheses about how inference is done and how concepts are noted. Because Genesis is the embodiment of such hypotheses, Genesis makes various kinds of heuristic inferences and uncovers conceptual content, just as we humans do when we understand stories.

Clark Glymour offers an analogy with epicycles to explain that good models not only should express natural behaviors, they should also exclude unnatural behaviors (2007). He argues that epicycles were not good models of planetary motion because you can use them to approximate any sort of motion to any degree of accuracy you want. There is, consequently, no opportunity for traditional falsifiability. The planets could move along the sides of squares and you could still explain what they are doing with epicycles. A better theory allows only the ellipses actually observed.

Does Genesis cover too much ground? Does Genesis do more than people can? With flexibility questions phrased that way, the notion of Turing completeness muddies the water, because given enough time and paper, a person, being a universal computer, could do anything any program can do. So we modify the question, asking not what people cannot do, but rather what people cannot do instinctively, or quickly, or normally.

The modified question brings us back to the computational-imperative principle. That is, nothing goes in unless Genesis needs it to do something people do instinctively, or quickly, or normally, so by construction, we avoid models that are so general they could explain behaviors different from or beyond those we humans exhibit. Of course, there could be some emergent behavior that would not be human, and that would falsify, but we have observed no such falsifying emergent behavior.

What about the other direction? Is there something people do that Genesis cannot do? Of course. Genesis is a emerging model of basic aspects of human story understanding, not a complete model of all human story understanding.

So we tend, by adherence to the doctrine of computational imperatives, to avoid models that are so general they can explain anything; and where we develop models that are narrow in scope, they are narrow for the uninteresting reason that we have much left to do.

We also do not introduce mechanisms that would only move us toward applications or enable statistical testing without shedding light on human story understanding. We believe that when human story processing is fully understood, there will be no shortage of exciting applications of that understanding.

1.9 What about other qualities?

What qualities, other than falsifiability, would determine whether Genesis or some other model embodies a better account of human story understanding? These are obvious: good models explain and predict; good models provide a unifying framework; good models are simple, honoring Occam's razor; good models are biologically and evolutionarily plausible; and good models support and benefit from Brain and Cognitive Science.

We are pleased that Genesis exhibits a range of story-understanding competences, offering behavioral explanations and predictions, on top of a substrate of hypothesized mechanisms that is sufficiently simple and small to be biologically and evolutionarily plausible.

1.10 What about engineering?

We certainly do engineering, because in developing Genesis, we are developing a prototype system with many capabilities such as those described in section 3. We like to think Genesis is analogous to the Wright Flyer of 1903.

1.11 What about evaluation?

If our purpose were to advance the state of the art in question-answering applications, we would of course find ways to compare Genesis with powerful question-answering applications such as Watson (Ferrucci *et al.*, 2010). We would also deploy mechanisms aimed at contest winning, whether or not those mechanisms plausibly shed light on human story understanding.

But our purpose is not immediate application. Our purpose is to devise and build a plausible account of human story understanding, showing not only how a story understanding system can answer questions, but also, as previously mentioned, describe conceptual content, summarize, compare and contrast, react with cultural biases, instruct, reason hypothetically, solve problems, and find useful precedents. When other systems exhibit such a range of capabilities, all resting on a common substrate, then it will be time to develop public test sets and make statistical comparisons.

1.12 On what prior research do we build?

In previous papers (2012a; 2012b), Winston introduced the Genesis Story Understanding System, emphasized methodological steps, articulated the Strong Story Hypothesis, and saluted the pioneering work of Roger Schank and his colleagues and students, documented in numerous articles and books Schank (1972); Schank and Abelson (1977); Schank and Riesbeck (1981); Schank (1991). In section 7, we explain in detail the contributions not only of Schank, but also of Marvin Minsky and many others who have contributed to story understanding research, with special attention to those whose work has influenced us in our development of Genesis.

1.13 Summary

Along with Berwick and Chomsky, we assume that we humans are the symbolic species because we humans have a unique operation, merge, that enables the construction of symbolic descriptions of properties, relations, actions, events, and constraining connections in an inner language. The ability to construct inner-language symbolic descriptions is what being symbolic means. That ability is an essential enabler of our uniquely powerful story understanding competence.

An inner story is a collection of usefully connected inner-language descriptions, which may be externalized to form an outer story expressed in an outer language, such as English, or in some other medium, such as video.

In our research, we follow methodological steps derived from those articulated by David Marr: we start with a specification of behavior; then we formulate computational problems; then we posit constraint exposing representations; then we build and test systems; and finally we articulate what has been learned. Our computational-imperative principle guides us toward falsifiable science and away from too-general explanations.

In the next section, we describe Genesis, whose development emerged from a desire to take steps toward an account of our human story understanding competence. Our purpose is to exhibit, in some detail, the representations and computations that we believe any such story-understanding system needs if it is to read text, absorb what it reads, make heuristic inferences, extract conceptual content, and exhibit various forms of humanlike understanding.

2 Genesis embodies steps toward an account of human story understanding

In this section, we explain how Genesis works via an explanation of the elements shown in figure 1. After exhibiting representative stories, we describe essential representations, comprised of classification hierarchies, case frames, and constraining connections. Then, we introduce common-sense rules and concept patterns, and we show how those common-sense rules and concept patterns enable Genesis to perform basic story understanding.

2.1 Genesis reads simple, concise stories

The following crime story, loosely based on Shakespeare's *Macbeth*, is representative of the stories on which we test our ideas:

> Maria is Fabiano's wife. Alegra is Michael's wife. Maria is evil and greedy. Vito is a godfather, and Fabiano is Vito's successor. Vito is an enemy of Luciano. Fabiano is brave. Fabiano defeats Luciano. Vito becomes happy because Fabiano defeats Luciano. Vito executes Luciano. Fabiano becomes chief of staff of Luciano. Vito rewarded Fabiano because Vito became happy. Fabiano is weak and vulnerable. Maria persuades Fabiano to want to become the godfather because Maria is greedy. Fabiano wants to become godfather because Maria persuaded Fabiano to want to become the godfather. Fabiano invites Vito to dinner. In order to murder Vito, Fabiano poisons Vito's food. Fabiano becomes godfather. Fabiano's murdering Vito leads to Michael's fleeing to Italy. Michael's fleeing to Italy leads

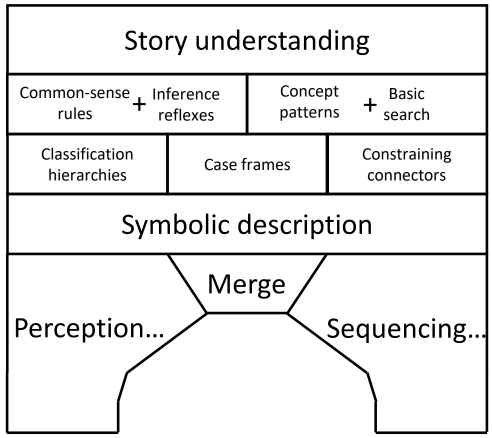

Figure 1: Genesis's story understanding rests on a small number of surprisingly simple representations and computations. At the representation level, Genesis requires classification hierarchies, case frames, and constraining connectors. At the computational level, much is done with common-sense rules via inference reflexes, described section 2.5, and concept recognition via basic search, described in section 2.9. All these are enabled by what appears to be our uniquely human, merge-enable keystone ability to build symbolic descriptions, along with other widely shared enablers.

> to Fabiano's murdering Alegra. Fabiano's murdering Vito leads to Maria's becoming distraught. Maria has bad dreams. Maria thinks God will punish her. Maria tries to pray. Maria kills herself. Michael returns from Italy and kills Fabiano.

The infrastructure and much of the knowledge developed to deal with conflict among people, as in *Macbeth* or the crime story, transferred over to other kinds of conflict, including, for example the Estonia–Russia cyber war of 2007:

> Estonia built Estonia's computer networks. Estonia insulted Russia because Estonia relocated a war memorial. Someone attacked Estonia's computer networks. The attack on Estonia's computer networks included the jamming of the web sites. The jamming of the web sites showed that someone did not respect Estonia. Estonia created a center to study computer security. Estonia believed other countries would support the center.

In both stories, harm causes anger: in one story because a person harms a person, and in the other story, because a country harms another country. Both situations are handled by a single common-sense rule like those described in section 2.5. Similarly, in both stories, harm leads to harm, in one story because people harm each other and in the other story, because countries harm each other. Both situations conform to the *Revenge* concept pattern described in section 2.9.

Other stories are adapted from psychological studies, such as the following murder story, based on a story used in a psychological study aimed at understanding cultural differences in how people explain violence (Morris and Peng, 1994).

> I am Chinese. America is a country. Lu is a student. Shan is a student. Lu comes from China. Shan comes from China. Lu studies in America. Shan studies in America. Lu fails his dissertation defense. Goertz is Lu's advisor. Lu dislikes Goertz. Lu had highest entrance exam score. Lu is a bachelor. Lu is lonely. Lu owns a gun. Lu practices shooting. Lu passes his second dissertation defense. Lu becomes a lab assistant because Lu fails to find a job. Shan graduates with Lu. Shan received national award. Faculty rejected Lu's appeal.

> Goertz angers Lu. Shan comes from a small Chinese village. Shan is married. Shan is social. Shan has friends. Shan angers Lu because Lu envies Shan. In order to kill Goertz, Lu shoots Goertz. In order to kill Shan, Lu shoots Shan. In order to kill Lu, Lu shoots Lu.

Finally, we used a brief synopsis of Shakespeare's *Macbeth* in our early work.

> Macbeth is a thane and Macduff is a thane. Lady Macbeth is evil and greedy. Duncan is the king, and Macbeth is Duncan's successor. Duncan is an enemy of Cawdor. Macduff is an enemy of Cawdor. Duncan is Macduff's friend. Macbeth defeated Cawdor. Duncan becomes happy because Macbeth defeated Cawdor. Witches had visions and danced. Macbeth talks with Witches. Witches make predictions. Witches astonish Macbeth. Macbeth becomes Thane of Cawdor. Duncan rewarded Macbeth because Duncan became happy. Macbeth wants to become king because Lady Macbeth persuaded Macbeth to want to become the king. Macbeth invites Duncan to dinner. Duncan goes to bed. Duncan's guards become drunk and sleep. Macbeth murders Duncan. Macbeth murders guards. Macbeth becomes king. Malcolm and Donalbain flee. Macbeth's murdering Duncan leads to Macduff's fleeing to England. Then, Macduff's fleeing to England leads to Macbeth's murdering Lady Macduff. Macbeth hallucinates at a dinner. Lady Macbeth says he hallucinates often. Everyone leaves. Macbeth's murdering Duncan leads to Lady Macbeth's becoming distraught. Lady Macbeth has bad dreams. Lady Macbeth thinks she has blood on her hands. Lady Macbeth kills herself. Birnam Wood is a forest. Birnam Wood goes to Dunsinane. Macduff's army attacks Macbeth's castle. Macduff curses Macbeth. Macbeth refuses to surrender. Macduff kills Macbeth.

We do not use Macbeth much in explanation because it is difficult to avoid thinking about whether our synopsis faithfully captures the nuances of Shakespeare's real play, instead of thinking about whether Genesis faithfully interprets the substance of the synopsis.

2.2 Genesis uses case frames extensively

Genesis uses Boris Katz's START system to translate simple English into a collection of descriptive entity-relation-entity triples (1997), which Genesis further processes into descriptions of story elements describing classifications, properties, relations, actions, and events.

Actions are expressed as case frames in the style of Charles Fillmore (1968). Then, depending on the action, there are various role players, such as the agent, object, co-agent, beneficiary, instrument, or conveyance. When the action involves motion, role players may include, for example, a source, destination, and direction.

Genesis uses entities, functions, relations, and sequences as a universal substrate for expressing story elements. An entity consists of a name along with a distinguishing index that ensures that two different entities with the same name are kept separate.

Functions are entities plus a subject slot filled by an entity or an entity subclass. Relations are functions plus an object slot filled by an entity or an entity subclass. Sequences are entities that hold either an ordered list or an unordered set of elements, each of which is an entity or an entity subclass.

Consider, for example, the sentence "A bird flew to a tree." When translated into a case frame, the action is *fly*, the bird is the *agent* and the tree is the *destination*.

When the case frame is expressed in the universal substrate of entities, functions, relations, and sequences, there are entities corresponding to the *bird* and the *tree*. The `tree` entity is the subject of a `to` function indicating a destination role. The `to` function is the sole element in a sequence holding a set of the role fillers. A `fly` relation connects the role-filler sequence to the subject, which by convention is taken to contain the agent role.

One way of displaying such a case frame follows. In section 2.7, we explain why we do not replace the *to* preposition with a destination-indicating symbol. Note that distinguishing indexes are not shown:

```
(relation fly
          (entity bird)
          (sequence roles (function to (entity tree))))
```

2.3 Genesis connects causes to consequents and means to actions

Genesis uses the same entity-function-relation-sequence apparatus to connect causing elements and caused elements, as in *a bird flew to a tree because a cat appeared.* In this example, there is just one causing element, the translation of *a cat appeared* and one element caused, the translation of *a bird flew to a tree*. All the causing elements are bundled together into a sequence, in this example expressing a one-element set. Then, the sequence of causing elements is tied to the element caused with a *cause* relation:

```
(relation cause
         (sequence conjunction (relation appear (entity cat) (sequence roles)))
          (relation fly
                    (entity bird)
                    (sequence roles (function to (entity tree)))))
```

Similarly, the entity-function-relation-sequence apparatus is used to connect means to actions. The following expresses the means specified in the sentence *In order to help Peter, Paul gave Peter advice:*

```
(relation means
          (sequence recipe
                    (relation give
                              (sequence roles (function object (entity advice))
                                              (function to (entity peter)))))
          (relation help
                    (entity peter)
                    (sequence roles (function object (entity paul)))))
```

Thus, stories are sequences of story elements, primarily represented as case frames and various kinds of connections that either appear explicitly in the English or that are inferred as described in section 2.5.

2.4 Genesis uses classification threads to capture class membership

Each entity and entity subclass also includes one or more classification sequences obtained by Genesis from WordNet (Fellbaum, 1998). The following, for example, are portions of two of the classification sequences obtained from WordNet for *hawk*:

```
thing entity physical-entity ... living-thing ... bird bird-of-prey hawk
thing entity physical-entity ... living-thing ... adult militarist hawk
```

2.5 Genesis uses common-sense rules to elaborate on what is written

Genesis uses common-sense rules to connect story elements. In our vernacular, whenever Genesis applies a common sense rule, we say that there has been an inference reflex.

Genesis uses actions and other story elements, together with inference reflexes, to build an *elaboration graph*, as shown in figure 2.

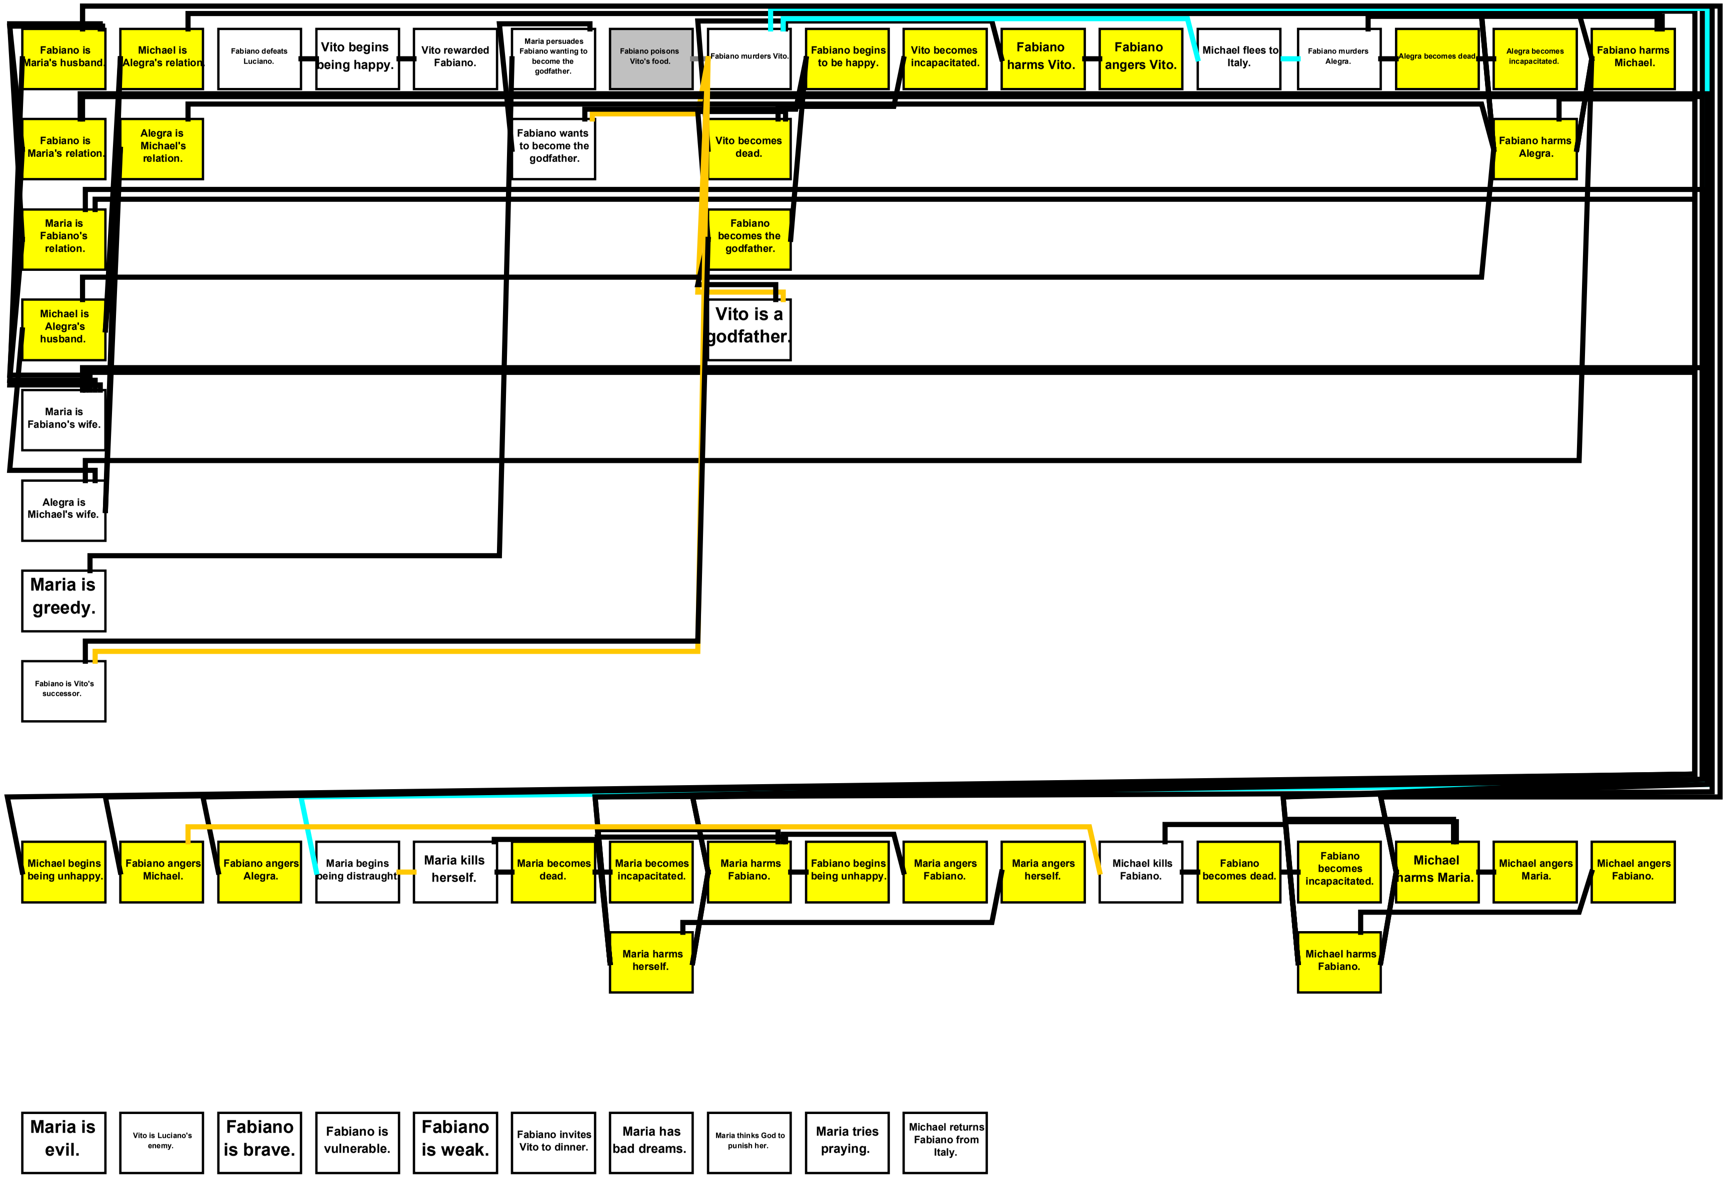

Figure 2: Genesis produces elaboration graphs as shown for a summary of the crime story. Elements in white are provided explicitly in the story. Elements in yellow are established by inference reflexes. Black connections represent deductions. Orange connections are inserted by explanation rules and cyan connections reflect sentences in the story that contain leads-to relations, both explained elsewhere in this section. Note that although the story is told as a sequence of elements, the inference reflexes form long-distance connections.

2.5.1 Deduction rules

We provide Genesis with deduction rules explicitly, expressing each in simple English, as in the following example:

```
If X kills Y, then Y becomes dead.
```

Here is the same deduction rule, translated from the English outer language into Genesis's inner language and expressed in the entity-function-relation-sequence substrate:

```
(relation cause
          (sequence conjunction
                    (relation kill
                              (entity x)
                              (sequence roles (function object (entity y)))))
          (function appear
                    (relation property
                              (entity y)
                            (sequence roles (function object (entity dead))))))
```

Figure 3 shows a portion of the elaboration graph of figure 2 centered on the application of a typical deduction rule.

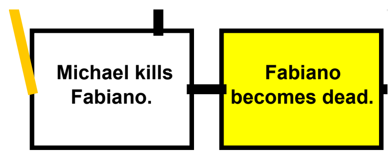

Figure 3: A deduction rule in action. The yellow box indicates the contents are a consequent.

2.5.2 Explanation rules

Whenever all the antecedents of a deduction rule appear in a story, Genesis asserts the consequent. Genesis uses deduction rules extensively, but if all Genesis had were always-true deduction rules, Genesis would seem quite stupid, because human thinking is not Aristotelian logic. We have found we need many common-sense rule types to model how humans digest stories.

For example, in reading a story, we humans seek explanations, and if none is offered, we assume connections that may hold, but not with sufficient regularity to be added by deduction rules. In the crime story in section 2.1, the story itself supplies no explicit reason why Fabiano murders Vito and no deduction rule supplies a reason. However, an explanation rule connects the murder to Fabiano's wanting to be godfather, Fabiano being Vito's successor, and Vito's being godfather.

Thus, Genesis does not assert the consequent of an explanation rule whenever the antecedents appear in a story; explanation rules make connections, but only if both the antecedents and consequent have already appeared and there is no known cause for the consequent.

Because Genesis uses explanation rules in an effort to find an explanation, given an observation, Genesis's use of explanation rules can be viewed as a kind of abduction. In our vernacular, however, abduction rules, explained on page 13, are specifications for more aggressive attempts at explanation that presume not only connections but also otherwise absent antecedents.

We express explanation rules in English using *may* as in the follow example:

```
If X is godfather and Y wants to be godfather and Y is X's successor,
then Y may murder X.
```

Another explanation rule connects anger to killing; fortunately, we do not always kill people who anger us, but it is a possibility:

```
If X angers Y, Y may kill X.
```

Figure 4 shows a portion of the elaboration graph of figure 2 centered on the application of a typical deduction rule.

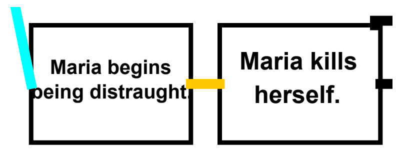

Figure 4: An explanation rule in action. Lacking any other explanation, an explanation rule ties Maria's suicide to becoming distraught. The yellow line indicates that the consequent is connected to the antecedent by an explanation rule.

Of course, other conventions would work as well as idioms to identify rule types, such as using explicit markers:

```
Explanation: If X angers Y, Y kills X.
```

Jonathan Gottschall powerfully supports the idea that we humans are explanation seekers in his seminal book, *The Story Telling Animal: How Stories Make us Human* (2012). He notes that when there is no explanation, we tend to make one up; he notes that because we are explanation seeking, manipulators can keep us in line by telling us appropriate stories.

2.5.3 Post hoc ergo propter hoc rules

A *post-hoc-ergo-propter-hoc rule*, also known as a *right-together rule*, is a kind of explanation rule that makes a connection only if the antecedent and consequent elements appear right together in a story. Such a connection is an error in logic, but perfectly natural in story reading. We express such rules using yet another an idiomatic expression:

```
X becomes Y; Z becomes angry.
```

Such a rule would make a connection if a story contained: "John became rich. George became angry." There would be no such connection if the story contained: "John became rich. It was a sunny day. Birds sang. George became angry."

Post-hoc-ergo-propter-hoc rules function much like what Eugene Charniak calls *demons* in his early thesis (1972).

We note in passing that when an author assumes a reader will engage a post-hoc-ergo-propter-hoc rule, that author is adhering to Grice's Maxim of Quantity (1989): the author leaves out the *because*, supplying no more information than is required.

2.5.4 Abduction rules

Some people consider murder to indicate insanity. We capture such thinking in an *abduction rule*, using a *must* idiom:

```
If X murders Y, then X must be insane.
```

Such a rule ensures that if there is a murder in a story, then the murder is a consequence of insanity. That is, if *John murders Peter* appears in a story, then the result is as if the story explicitly included *John murders Peter because John is insane*.

Note that abduction rules can specify antecedent actions, not just antecedent characteristics:

```
If X hates Y, then Y must anger X.
```

Note also that when Genesis uses an abduction rule, Genesis explains a consequent by presuming both an absent antecedent and a causal link. Thus, Genesis's abduction rules presume more than Genesis's explanation rules and post-hoc-ergo-propter-hoc rules, which explain consequents by presuming only a causal link to one or more antecedents already in place.

2.5.5 Presumption rules

A *presumption rule*, like an abduction rule, assumes a particular cause, indicated by a *can be* idiom:

```
X can be greedy because X is evil.
```

With such a rule in place, if *John is greedy* appears in a story, and there is no explicit cause or cause put in place by a rule, then the result is as if the story explicitly included *John is greedy because John is evil*.

2.5.6 Enablement rules

An *enablement rule* supplies essential prerequisites to an action. Enablers appear in *enables* idioms:

```
X's having a knife enables X's stabbing Y.
```

Whenever a stabbing occurs, Genesis concludes that the stabbing person must have a knife, and that the having and stabbing are connected by an *enables* relation.

2.5.7 Censor rules

Another kind of rule, a *censor rule*, prevents inappropriate inference, as when a deduction rule might otherwise make a dead person unhappy. A *cannot* idiom identifies this kind of rule:

```
If X becomes dead, X cannot become unhappy.
```

Thus, if the antecedent of a censor rule is present, the consequent cannot be asserted by any other rule.

2.6 Inference reflexes

We have, so far, identified six rule types that establish connections, loosely considered kinds of cause, and one rule type that prevents connection. Each is expressed in an idiom, with the exact form of the idiom jointly constrained by what the front-end START parser can handle and by a desire to have all knowledge in human-readable form. Here is how they work together whenever a new story element is added to a story:

1. When a new element is an antecedent to a *deduction rule* and other antecedents of that rule are already present in the story, Genesis uses the deduction rule to assert a conclusion and constructs a deduction connection between the antecedents and the new conclusion.

2. When a new element has no cause and the new element is the consequent of a *explanation rule*, and the antecedents of the explanation rule are already present in the story, Genesis uses the explanation rule to construct an explanation connection between the antecedents and the conclusion.

3. When a new element still has no cause and the new element is the consequent of a *post-hoc-ergo-propter-hoc rule*, and the antecedent of that rule is present and lies immediately before the new element in the story, then Genesis uses the post-hoc-ergo-propter-hoc rule to construct a proximity connection between the antecedent and the consequent.
4. When a new element still has no cause and the new element is the consequent of an *abduction rule*, then Genesis uses the abduction rule to enter the antecedent of the rule into the story and to construct an abduction connection between the newly entered antecedent and the consequent.
5. When a new element still has no cause and the new element is the consequent of a *presumption rule*, then Genesis uses the presumption rule to enter the antecedent of the rule into the story and to construct a presumption connection between the newly entered antecedent and the consequent.
6. When a new element is the consequent of an *enablement rule*, enter the antecedents that must be true for the consequent to occur and construct an enablement connection between the antecedents and the consequent.
7. Finally, if any deduction, abduction, presumption, or enablement rule attempts to enter a story element that is the consequent of a *censor rule* and the antecedents of the censor rule are already present in the story, Genesis terminates its attempt to enter the forbidden story element.

We discovered each such rule type when working to model human reaction to particular stories, not through a design exercise disconnected from any specific case. Take away any rule type and Genesis could not understand some story properly. Accordingly, each rule type constitutes a *computational imperative*.

Each rule does its work the moment it can, and because each application of a rule is a sort of knee jerk in response to circumstance, we call each application an *inference reflex*.

> **Inference reflex:** The automatic addition to an inner story of an element or connection between elements using a common-sense rule.

2.7 Common-sense rules retain prepositional markers

Note that if all we care about is how a story element matches a rule's antecedents or consequent, and if a rule is described with the same case-marking prepositions that appear in a story, and if a rule contains antecedents that constrain what kind of things are matched, then we can defer role interpretation from rule reading time to rule use time.

Consider, for example, *Peter killed Paul with Mary*, and *Peter killed Paul with a wrench*. In the first sentence, Mary is a co-agent; in the second, the wrench is an instrument. The preposition *with* can introduce either, but the following rules make the correct inferences nevertheless:

```
If W is a living-thing and X kills Y with W, then W is an accomplice.
If W is a artifact and X kills Y with W, then W is a weapon.
```

Mary is a person, and according to WordNet, a person is a `living-thing`; a wrench is a tool and a tool is an `artifact`. Accordingly, the first rule makes only Mary an accomplice and the second rule makes only the wrench a weapon. The correct action can be sorted out by matching at inference-reflex time because the common-sense rules specify what should match.

2.8 Explicit connections also contribute to basic understanding

Of course, a story may itself exhibit causal connections, as in an *explicit cause* statement:

```
Vito begins being happy because Fabiano defeated Luciano.
```

Alternatively, a connection may involve a chain of causes, with only the first and final elements mentioned in a *leads to* statement:

```
Fabiano's murdering Vito leads to Michael's fleeing to Italy.
```

Figure 5 show a portion of the elaboration graph of figure 2 centered on where an explicit *leads to* connection appears.

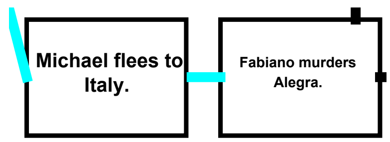

Figure 5: Blue connector marks the use of a leads-to expression.

Some leads-to statements come with an explicit indication that you will never understand the details. Such an *unknowable leads-to* statement is expressed using a *strangely* idiom:

```
Strangely, Fabiano's murdering Vito leads to Maria's beginning to be distraught.
```

The use of a semicolon forces a post-hoc-ergo-propter-hoc connection even in the absence of a right-together rule:

```
Michael kills Fabiano; Michael is happy.
```

Still another connection expresses how an event occurs. We call these *means* expressions; they appear in *in order to* idioms:

```
In order to murder Vito, Fabiano poisoned Vito.
```

2.9 Genesis reflects on its reading, looking for concepts

Once Genesis builds the elaboration graph, Genesis looks for instances of concepts (Nackoul, 2010) using concept patterns that specify elements and connections among them, building on the pioneering work of Wendy Lehnert (1981).

> **Concept recognition:** The affirmation that an inner story contains the elements and connections that appear in a concept pattern.

The following, for example, is a concept pattern for *Revenge*. The leads-to relation indicates that there is a sequence of causal connections between the harming actions:

```
Start description of "Revenge".
X is an entity.
Y is an entity.
X's harming Y leads to Y's harming X.
X is not equal to Y.
```

In figure 6, Genesis notes a *Revenge* pattern because Genesis successfully searches for a sequence of causal connections between Fabiano's harming Michael and Michael's harming Fabiano.

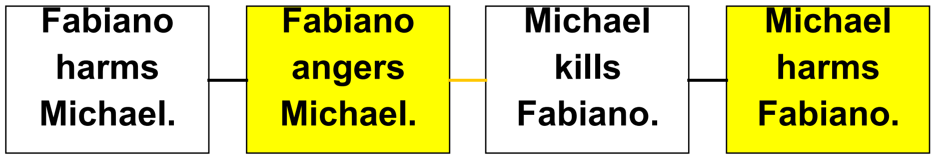

Figure 6: Genesis finds concept patterns by searching the elaboration graph. Here, Genesis highlights *Revenge* elements in the elaboration graph. The inspector panel provides a close-up view.

Some concept patterns specify more elaborate connections, such as the following for *Pyrrhic victory*:

```
Start description of "Pyrrhic victory".
X is an entity.
Y is an entity.
A is an action.
X's performing A leads to X's becoming happy.
X's performing A leads to Y's incapacitating X.
Y incapacitates X after X becomes happy.
```

In figure 7, Genesis notes a *Pyrrhic victory* pattern because Fabiano's wanting to be godfather leads not only to becoming happy, but also leads to being incapacitated by someone later.

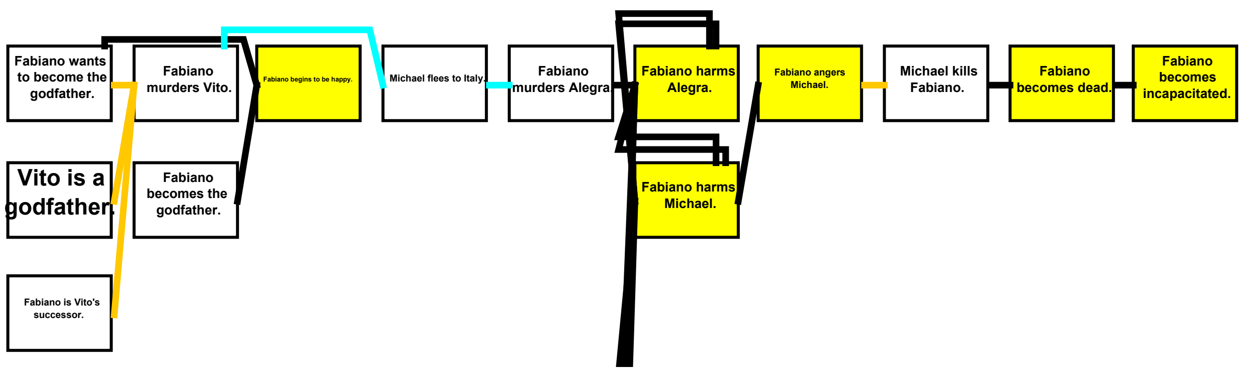

Figure 7: Genesis extracts the Pyrrhic victory elements from the full elaboration graph. Black connectors mark the use of deduction rules; yellow connections mark the use of explanation rules. Blue connections indicate that the story contains leads-to expressions.

Most concept patterns, but not all, are like *Revenge* and *Pyrrhic victory* in that they include leads-to relations. Thus, concept identification generally requires search, which takes concept recognition beyond the reach of common-sense rules as ordinarily used.

An optional concept-pattern element, the *sometimes* element, specifies that an entity may or may not be present in a story, but if it is, it becomes part of the recognized concept. The following, for example, is another version of *Revenge*; there may or may not be hating between the participants:

```
Start description of "Revenge".
X is an entity.
Y is an entity.
X's harming Y leads to Y's harming X.
Sometimes X hates Y.
Sometimes Y hates X.
```

Another optional concept pattern element, the *consequently* element, specifies an entity that is to be inserted back into a story as a by-product of noting a concept is present. The following emerged in work with Native American Crow creation myths (Yarlott, 2014):

```
Start description of "Violated belief - Medicine Man".
X is a person.
Y is an thing.
Z is an thing.
X transforms Y into Z.
Consequently, X has strong medicine because X transforms Y into Z.
```

2.10 Concept patterns enable abstraction

Note that *Revenge* is an abstraction identified with harming events. The particular kind of harming event is unimportant; it may involve a mild insult or a vicious killing. As long as two harming events are connected, with the harms going in opposite directions, there is *Revenge* in a story. The word *revenge* or a synonym need not appear, so no system that looks only at words can reliably identify *Revenge*.

2.11 We are our own knowledge engineers

Soon after rule-based expert systems were developed in the 1980s, those who built them began to call the enterprise of rule development *knowledge engineering*. Two key heuristics soon emerged (Winston, 1992):

- Look at specific cases. Rules come to mind when looking at specific cases that do not come to mind when thinking in general terms.
- Ask why situations that seem superficially the same are not the same. Often the answers suggest analysis requires additional vocabulary.

All the rules and concepts we have composed come from looking what is needed to handle specific cases. Sometimes a specific case even forces the introduction of a new, unanticipated kind of rule. We did not know we would need explanation rules, for example, until early work on our Macbeth summary left us puzzled about why Genesis did not see *Revenge*. There was no complete harm-leads-to-harm path until we added a new explanation rule, *If X angers Y, Y may kill X*.

Often, new vocabulary emerges as we examine why two situations we know to be different seem to Genesis to be the same. Genesis saw an excess of *Pyrrhic victory* concepts until we replaced *Y harms X after X becomes happy* with *Y incapacitates X after X becomes happy*. There is no Pyrrhic victory, after all, if the original consequence is more positive than the downstream consequence is negative.

Thus, in composing rules and concepts, we use the key heuristics of knowledge engineering, so it is reasonable to say we are acting as our own knowledge engineers.

In section 5, we discuss how Genesis might acquire rules and concepts by asking questions, by connection to large repositories of knowledge, and by self discovery.

2.12 Summary

Genesis's essential representational foundation consists primarily of classification threads to capture classification information, case frames to express actions, and various kinds causal connections to establish constraint.

Genesis uses five kinds of explicit causal connections, six kinds of common-sense rules to make causal connections, and censor rules that prevent inappropriate use of deduction, abduction, presumption, and enablement rules. Genesis uses concept patterns to specify entities that must be present and entities that must be causally connected. The concept patterns may exhibit two kinds of optional elements.

The explicit elements translated from a story into Genesis's inner language, augmented by elements produced by various kinds of common-sense rules and concept patterns, constitute a Genesis inner story.

Over time, more representational, common-sense rule and concept pattern types will be discovered, but what has already been demonstrated suggests that the number of types needed in an account of human story understanding is plausibly small.

3 Genesis's simple substrate supports many competences

In this section, we list some of the competences enabled by Genesis's small number of explicit connection types, common-sense rule types, and concept pattern types. Our purpose is to argue that a simple, plausibly evolvable foundation, suffices to support myriad competences such as those shown in figure 8. In section 7.2, we compare our work on these competences to that of others.

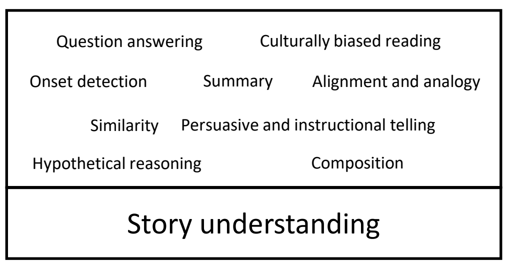

Figure 8: Basic story understanding, enabled by a small set of representational capabilities, common-sense rule types, explicit connection types, and concept pattern types, serves as a foundation on which we have built all the story-understanding competences described in this section.

Reuse of knowledge, of course, is an important property of any theory of story understanding because without reuse there would be no possibility of education and our evolving model would have a show-stopping flaw.

Accordingly, as the competences described in this section were developed, we were pleased to note that there was much reuse. Killing makes the victim become dead whether in a fairy tale or in a Shakespearean play. *Revenge* is the same concept whether involving people in a Shakespearean play or countries in a cyber war. Thus, as Genesis moves from one story to another, Genesis reuses a great deal of already recorded knowledge.

3.1 Aspects of many competences have been demonstrated

Here we describe implemented models of aspects of many story competences. We are excited by these illustrations, but we also understand that they are just illustrations of what can be done with stories written by us for Genesis, a step toward demonstrations in which large numbers of stories written by people for people are processed, with statistical analyses of what works and what fails. In section 5 we discuss Genesis's limitations in more detail.

3.2 Genesis answers basic questions about why and when

As shown in figure 9, Genesis answers questions on two levels, by reciting elaboration graph elements connected to the target event, and by noting how the target event is embedded in concepts.

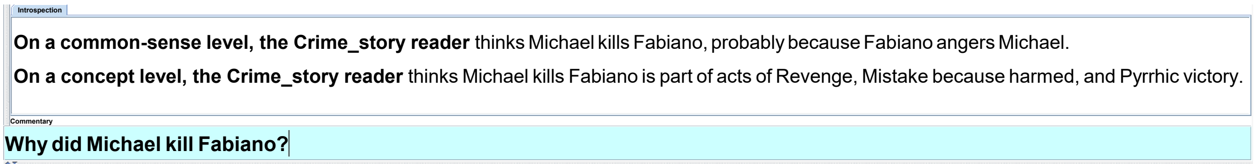

Figure 9: Genesis answers a user question, "Why did Michael kill Fabiano," on the common-sense level and the concept level.

3.3 Genesis models personality traits

Genesis notes what various sorts of people do, which in turn enables Genesis to infer personality traits on the basis of what people do, which enables Genesis to use personality traits to explain acts (Song, 2012).

Genesis notes early in one version of the crime story that Michael assaults someone, an act Genesis has recorded as indicative of vicious people, leading Genesis to consider Michael to be vicious. Then, whenever Michael is involved in an action, common-sense rules associated with viciousness are added to those generally used.

3.4 Genesis notes concept onsets, anticipates trouble

Concepts generally involve leads-to relations. Noting the first part of a leads-to relation provides early warning of possible evolutions. As shown in figure 10, the potential for four concepts is noted in the course of reading the crime story.

Figure 10: Genesis notes the onset of four possible concepts midstream in the crime story. *Revenge* is especially prominent as there are so many instances of *harm*, the first part of the leads-to relation in the *Revenge* concept.

3.5 Genesis reads stories with controllable allegiances and cultural biases

Genesis's interpretation may shift dramatically with a small shift in what a story contains. In the example shown in figure 11, based on the 2007 cyber war between Estonia and Russia, Genesis views the alleged actions of the Russians as *Aggression of a bully* because the story includes *Estonia is my ally*. Genesis views the same actions, as shown in figure 12, as *Teaching a lesson* because the story includes *Russia is my ally* (Winston, 2012b).

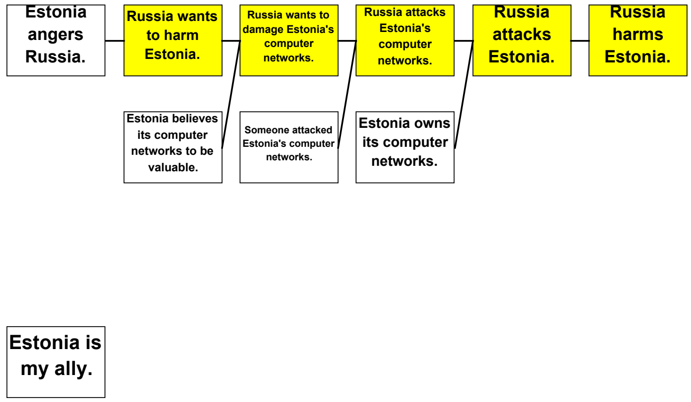

Figure 11: Genesis views the 2007 cyber war between Estonia and Russia from the perspective of an ally of Estonia. The result: *Aggression of a Bully*.

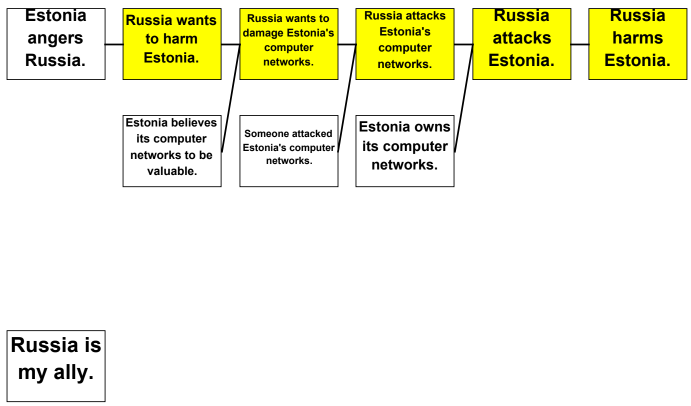

Figure 12: Same cyber war, but from the perspective of an ally of an ally of Russia. The result: *Teaching a Lesson*.

3.6 Genesis aligns similar stories for analogical reasoning

Genesis aligns stories, in preparation for analogical reasoning, using the Needleman-Wunch algorithm borrowed from molecular biology. In figure 13, Genesis finds clear parallels between the onset of the Arab-Israeli war and the Tet Offensive in the Vietnam war. In both cases, intelligence noted mobilization, intelligence determined that the attackers would lose, intelligence determined that the attackers knew they would lose, intelligence concluded there would be no attack, whereupon the attackers promptly attacked. Retrospectively, there were political rather than military motives.

Such alignments do not make certain predictions, but they can suggest how precedents may apply to current events, potentially stopping the kind of oversight blunders that are common in the fog of war (Fay, 2012).

3.7 Genesis models question-driven story augmentation and problem solving

After reading a story, a question may stimulate further analysis and expose new conclusions. The example here is from a Chinese-American story understanding demonstration (Morris and Peng, 1994; Awad, 2013) using the murder story shown in section 2.1.

In the story, Lu, a Chinese student, studying in America, murders a professor and another student. In their study, Morris and Peng asked subjects to weigh the importance of various factors contributing

Story	Element 0	Element 1	Element 2	Element 3
Story B Score: 6.0	Israelis know that Egyptians are preparing to attack them.	Egyptians are preparing to attack Israelis.	Israelis defeat Egyptians.	Israelis know to defeat Egyptians.
Story A Score: 6.0	USA knows that Viet Cong is preparing to attack it.	Viet Cong is preparing to attack USA.	USA defeats Viet Cong.	USA knows to defeat Viet Cong.

Element 4	Element 5	Element 6	Element 7	Element 8
Israelis know that Egyptians know they defeat Egyptians.	Egyptians know that Israelis defeat them.	Egyptians don't attack Israelis.	Israelis believe that Egyptians don't attack them.	Gapfilled, Egyptians attack Israelis.
Gapfilled, USA knows Viet Cong knows it defeats Viet Cong.	Gapfilled, Viet Cong knows USA defeats it.	Viet Cong doesn't attack USA.	USA believes Viet Cong doesn't attack it.	Viet Cong attacks USA.

Figure 13: Genesis aligns elements in two wars and fills gaps in each using the other.

to Lu's murders. Some of the subjects were American and some were Chinese; all were graduate students studying at the University of Michigan. The experiment produced the following averages:

Possibly contributing factor, 1 = not a cause and 4 = a major cause	American average	Chinese average
American movies and television glorify violent revenge tactics.	1.5	3.6
America's extremely individualistic, selfish values corrupt foreign students.	1.2	2.5
Lu had chronic personality problems.	4.2	2.4
Lu was mentally imbalanced because his life consisted only of work, without other activities which relieve stress.	4.5	1.8

Evidently, Chinese students attributed about twice the weight as Americans to situational causes and half the weight to personal dispositional causes.

In our demonstration, we focused not on weight but on connection. We asked what it would take to have Genesis connect the murders to the situational causes for a Chinese reader but not for an American reader.

Genesis, modeling a Chinese reader, does not at first see any connection between Lu's killing of Shan and America's glorification of violence because there is nothing about America glorifying violence in the story.

Then, as shown in figure 14, having been asked a question, Genesis recalls that the question's antecedent is something that the reader believes. That leads to Genesis's adding that recalled belief to the story. That, in turn, leads to connecting the inserted belief, America glorifies violence, to the murder, as shown in figure 15. Genesis then affirms that Lu killed Shan because America glorifies violence.

Another version of Genesis, modeling an American reader, recalls no such belief, so fails to insert *America glorifies violence* into the story. Thus, there can be no connection of antecedent to consequent. This time, Genesis denies that Lu killed Shan because America is individualistic.

3.8 Genesis develops summaries around conceptual content

Because Genesis recognizes conceptual content, Genesis can construct intelligent summaries by ignoring all the story elements that are not connected with a central concept.

In the following example, limiting the telling of the Lu murder story to those elements connected to the central murderous influence concept compresses the summary provided into a shorter summary by about 3:1 (Winston, 2015).

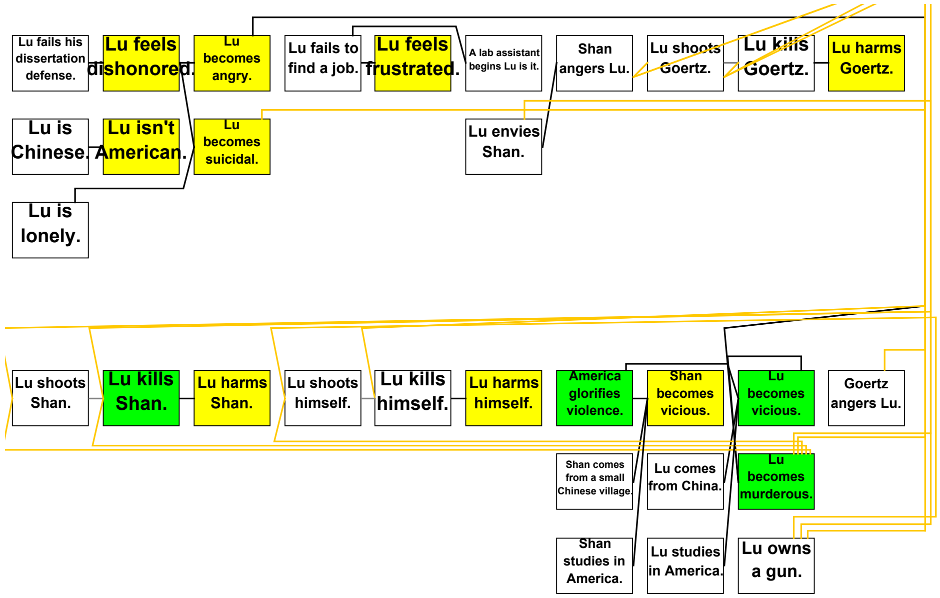

Figure 14: Genesis interprets the Shan murder. The basic interpretation does not connect the murder with America's glorification of violence until a question is asked and a belief is inserted by the Genesis version that models a Chinese reader. Then, Genesis, modeling a Chinese reader, notes a *Murderous influence* concept pattern connecting the glorification of violence to the murder.

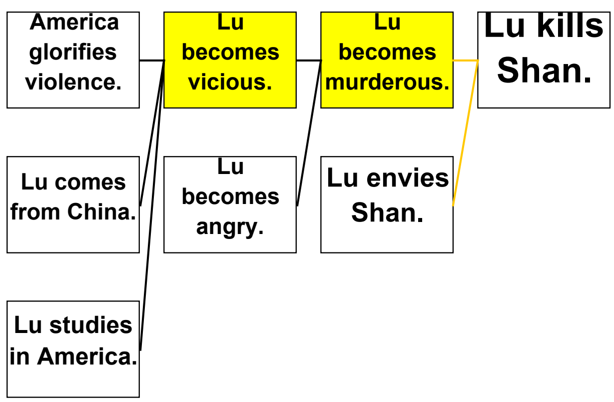

Figure 15: Genesis connects *America glorifies violence* to Lu kills Shan. Note that an explanation rule connects the killing to *Lu becomes murderous* and *Lu envies Shan*.

> The story is about Murderous influence. Lu comes from China. Lu studies in America. America glorifies violence. Lu fails his dissertation defense. Lu kills Goertz, probably because Lu becomes murderous, and Goertz angers Lu. Lu kills Shan, probably because Lu becomes murderous, and Lu envies Shan. Lu kills himself, probably because Lu becomes murderous. Story contains 42 elements, summary 13, or 31.0%.

3.9 Genesis calculates similarity using concepts

Genesis judges similarity in multiple ways. One way is by using word vectors; another is by using vectors whose components are concept counts. Using concept vectors enables Genesis to see similarities not evident in the words. The following two-sentence stories illustrate. All involve different actors; all involve *Revenge* because harm leads to harm; none uses the word *revenge*.

Story 1: The pig ate the dog's food. The dog bit the pig.
Story 2: John insulted Mary. Mary yelled at John.
Story 3: The paper criticized the party. The party threatened the paper.

The comparisons shown in figure 16 are on pairs of short descriptions of conflicts (Krakauer, 2012).

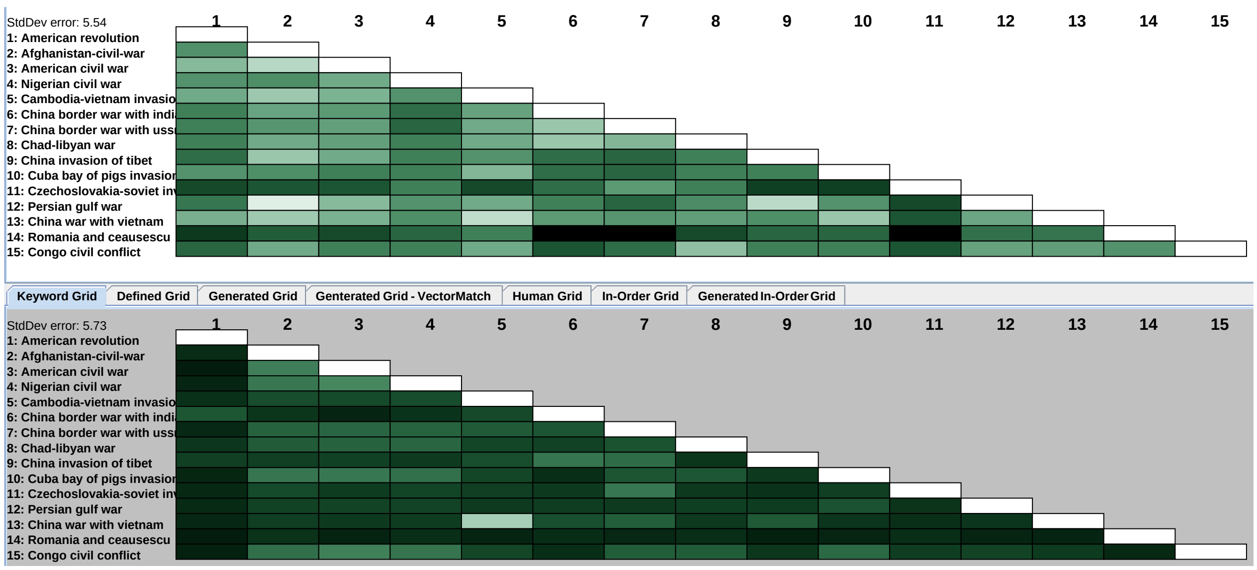

Figure 16: Genesis performs concept-based similarity measurements. Concept-based measurements are shown above and word-based similarity measurements below. White means most similar. As expected, the similarity conclusions reached using concepts are invisible to approaches using only words.

3.10 Genesis explains using a reader model

Using a model of what a story reader knows, Genesis can tailor telling to cover gaps in the reader's knowledge by simple spoon feeding, by more elaborate explanation, or by helpfully supplying principles (Sayan, 2014). In figure 17, Genesis supplies principles to a reader that knows very little in the beginning, but is taught that, for example, you become incapacitated if you become dead.

3.11 Genesis composes new stories

Many human authors say that once they have created elaborate character sketches, and place the characters in an initial situation, stories seem to write themselves. Presumably the sketches and the situation call to mind fragments from a story library, which the author then weaves together to compose a new story, thus exhibiting an aspect of creativity.

Matthew Fay captured that character-driven, fragment-assembly authoring idea in his story composition system (2014). Fay's system reused elements from Shakespearean tragedies and war stories:

Figure 17: Genesis uses a reader model to determine what and how much to say in retelling Crime Story. Here, Genesis says a lot, because Genesis's model of the reader suggests that the reader does not know much. Conclusions are in red text; reasons are in blue text.

> Greinia and Astilir
>
> Greinia is a country. Astalir is a city in Greinia. Angelina is from Astalir. Angelina is from Greinia. Angelina becomes the Queen of Greinia. Angelina rules Greinia. Malcolm is Angelina's successor. Malcolm kills Angelina. Malcolm becomes King of Greinia. Astalir rebels. Astalir attacks Greinia. Malcolm scares away Astalir's forces. Astalir attacks Malcolm. Astalir kills Malcolm. Astalir defeats Greinia. Astalir becomes independent.

Fay's system retells Hansel and Gretel with the Gretel character removed. Too bad for Hansel:

> Hansel without Gretel
>
> Mother dislikes Hansel. Mother wants to kill Hansel. Mother convinces Father to abandon Hansel in the woods. Father abandons Hansel in the woods. Hansel becomes hungry. Hansel finds the house made of candy. The witch lives in the house made of candy. The witch enslaves Hansel. The witch becomes hungry. The witch wants to eat Hansel. The witch fattens up Hansel. The witch pushes Hansel into the oven. The witch kills Hansel. The witch eats Hansel. Father discovers Hansel died. Father kills himself.

3.12 Genesis reasons about who knows what

Genesis infers what various characters know based on who is present and paying attention. An example, based on Les Mis⊠rables, demonstrates Genesis's who-knows-what ability (Noss, 2017):

> Inspector Javert is a policeman. Jean Valjean commits a crime. Then, Jean Valjean repents. Jean Valjean becomes a good person.

In the story, the perspectives of the policeman and criminal diverge because Javert is presumed to be absent when Jean Valjean repents and becomes a good person.

Using who-knows-what knowledge to construct character-specific elaboration graphs, Genesis answers questions about beliefs, retells stories from various character's point of view, and explains misunderstandings that arise between characters with different information or different biases. In figure 18, for example, Genesis answers a question about what Javert believes.

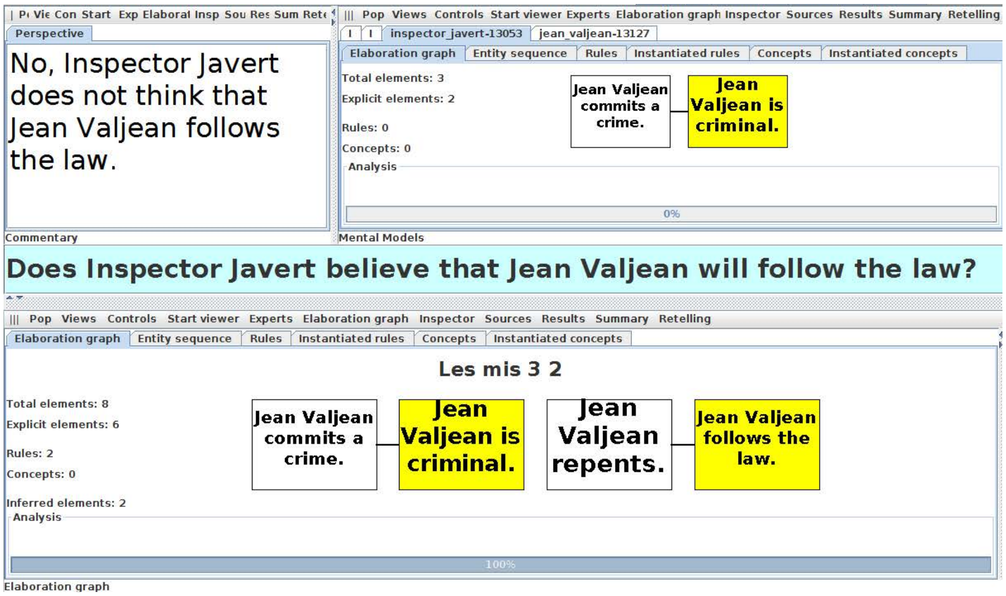

Figure 18: Genesis answers comprehension questions about characters' perspectives. Here, Genesis observes that the policeman will not believe that this criminal follows the law because the policeman was not present when the criminal repented.

Genesis can also compare two characters' perspectives, attributing differences in interpretation to differences in what is observed. The following records an exchange between a human questioner and the Genesis:

> Why does Jean Valjean disagree with Inspector Javert?
> Inspector Javert and Jean Valjean disagree about "Jean Valjean is criminal".
>
> Why does Jean Valjean think that Jean Valjean isn't criminal?
> Jean Valjean infers that Jean Valjean isn't criminal because [he] repents.
>
> Why did Inspector Javert think that Jean Valjean is criminal?
> Inspector Javert infers that Jean Valjean is criminal because [he] commits a crime.

We believe that Genesis's who-knows-what ability sheds light on our human ability to reason about what others know and believe. Genesis's who-knows-what ability captures aspects of common sense (being within earshot, being unconscious or distracted, speaking over the phone or in another language), provides tools to aid in diplomacy and education (pinpointing differences in knowledge and experience), and suggests computational explanations of various psychological disorders (defects in mechanisms that enable understanding what others think).

3.13 Genesis persuades

Similarly, Genesis can tailor what is said to shape reader opinion. In figure 19, for example, sentences that involve actions associated with likability are emphasized, while those associated with unlikability are deleted, so as to make the Woodcutter look good, and everyone else look bad, in Genesis's version of *Hansel and Gretel* (Sayan, 2014).

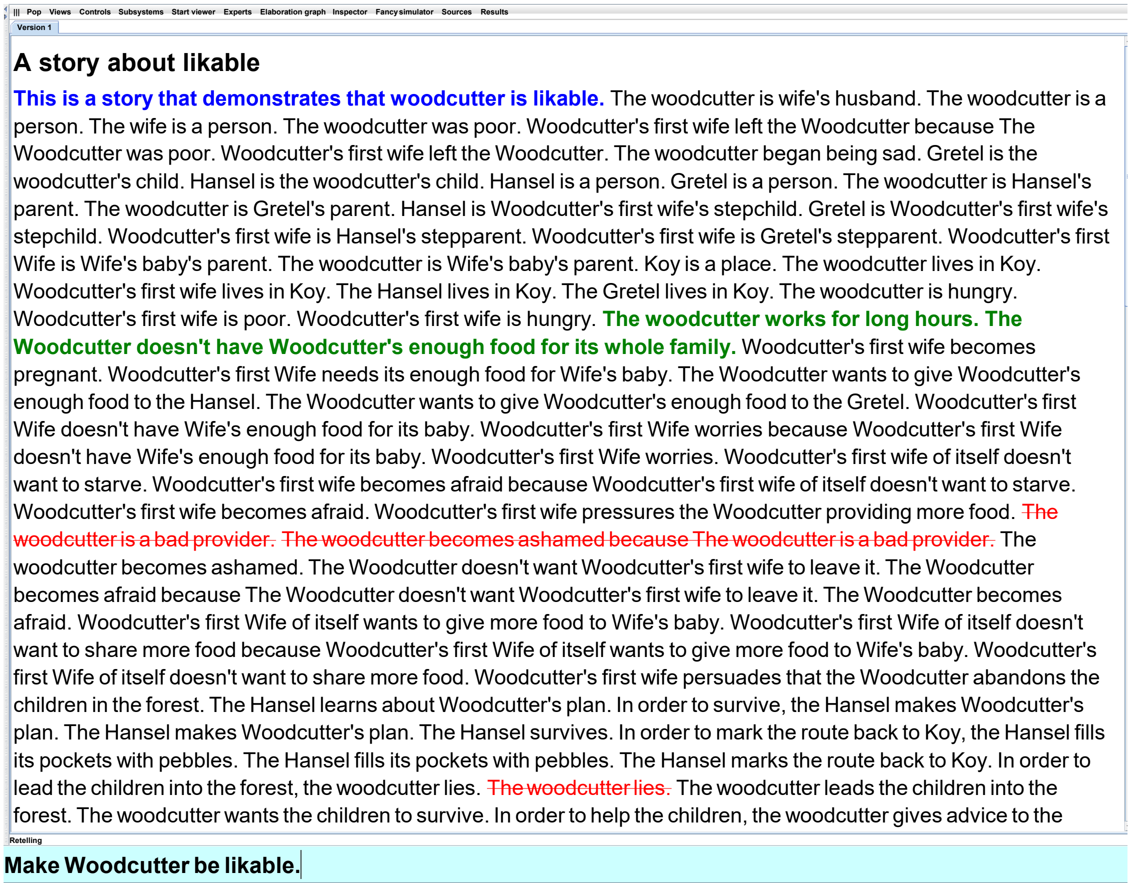

Figure 19: Genesis uses a reader model to determine what and how much to say so as to shape the reader's opinion in this first part of a retelling of Hansel and Gretel. The good is in bold green text; the bad is in red text and struck out.

3.14 Genesis reasons about hypothetical possibility

Hypothetical reasoning is another of Minsky's suitcase labels. It can refer to a variety of abilities, some of which we have begun to model using story understanding as a substrate.

When asked a hypothetical question about a story, Genesis reinterprets the story with indicated additions or deletions. Suppose, for example, that you ask Genesis for its analysis of the following story:

> Start story titled "Lover brandishes a knife". George, Alex, and Martha are persons. George is Martha's spouse. Alex is Martha's lover. Alex and George despise each other. George encounters Alex and Martha at a bar. George yells at Alex. Alex brandishes a knife. Alex becomes dead because George shoots Alex. The end.

It certainly looks like a case of self defense and Genesis reaches that conclusion, as shown in figure 20. But what should be concluded if Alex did not brandish a knife? To answer, Genesis removes the knife-brandishing element from the story, re-analyzes the story, and arrives at the conclusion shown in figure 21.

Genesis also produces a summary describing the differences at both a fine-grained story-element level and at an abstract, conceptual level:

```
From an event-based perspective, note the following changes:
It's no longer the case that Alex intends harming someone.
From a thematic perspective, the following concepts disappear:
Alex's self defense
... and the following concepts are introduced:
George's spiteful violence
```

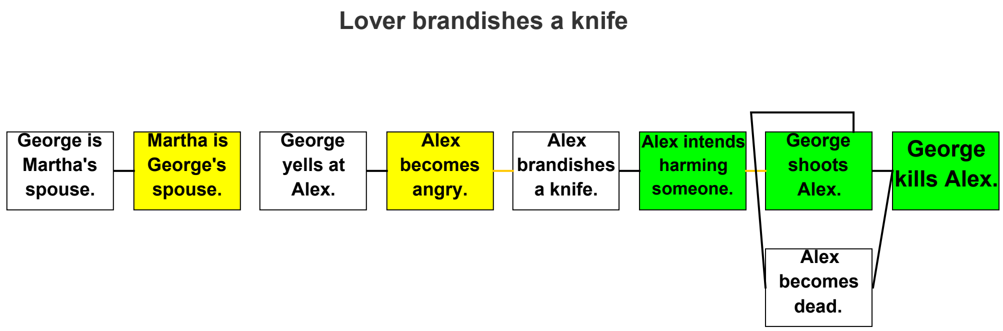

Figure 20: In the knife-brandishing version of a story raising a legal question, the elaboration graph indicates knife the brandishing is connected to the killing, suggesting *Self defense*.

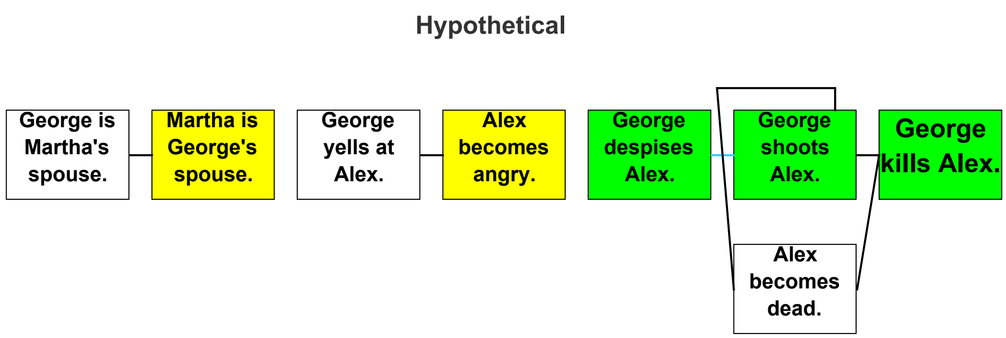

Figure 21: In the hypothetical version, with the brandishing removed, Genesis presumes the explanation for killing has to do with despising, suggesting guilt. At the concept level, *Self defense* becomes *Spiteful vengeance*.

We envision future systems for analysts who want to ask what-if questions about law, policy, or diplomatic intervention; such systems would be analogous to today's systems for financial analysts who use spreadsheets to ask what-if questions about best-case and worst-case scenarios.

3.15 Genesis's common sense rules add new elements and explanations

In the examples, the number of explicit elements and the ratio of explicit elements to deductions and explanations depends on how the experimenter chooses to tell the story. Here are representative examples:

Story	Explicit elements	Deductions	Explanations
Lover	14	4	2
Estonia/Russia	18	29	0
Les Mis	26	8	1
Crime	35	32	3
Lu Murder before question	41	8	1
Lu Murder after question	42	11	6
Crow creation myth	55	3	3
Macbeth	88	36	4
Hansel & Gretel	210	28	0

4 We view story understanding as foundational

Many researchers have come to see story understanding as an important subject. Most have approached story understanding as a challenging problem whose solution requires all the reasoning, planning, representation, and architectural apparatus that Cognitive AI has to offer.

We think it reasonable to take the complementary view: because we believe story understanding is the *sine qua non* of human intelligence for the reasons given in section 1, we believe story understanding is foundational. We think that recipe following is a kind of story understanding, that problem solving is a kind of recipe following, and that recipe construction is a kind of authoring.

4.1 Story understanding supports solving story problems

According, we built a problem solver, much like those of Pat Langley (Langley *et al.*, 2013, 2014; Bai *et al.*, 2015; Langley *et al.*, 2016), except that all of the knowledge is expressed in English in what we think of as micro stories (Winston, In preparation; Yang and Winston, In preparation).

Here, for example, are some of the problem solving micro stories used in working with the Shan murder example in section 3.1:

```
If the question is "Did xx cause yy".
Intention: Establish that a path leads from xx to yy.
The end.

If the intention is "Establish that a path leads from xx to yy".
Condition: Verify that xx is in the story.
Condition: Verify that yy is in the story.
Method: Search for a path from xx to yy.
The end.

If the condition is "Verify that xx is in the story".
Check: Look for xx in the story.
Solve: Do I believe xx.
The end.

If the method is "Search for a path from xx to yy".
Execute: Call "findPath" with xx with yy.
The end.
```

The final microstory specifies a call to a `findPath` method, a method we think of as a just-do-it method, one that lies outside the problem solver and is not subject to inquiry via the story that Genesis tells itself as its problem solver does its work.

If asked how you would drink from a glass, you would likely say you reach out, close your fingers, lift, move the glass to your mouth, tilt. But if asked how you close your fingers, you would have to say you just will it and it happens. You have grounded out in a just-do-it method like those used under the Genesis story-based problem solver.

4.2 Story understanding supports solving robot problems

What works for asking questions about a story works for robot problem solving as well. Here, for example, are some of the problem solving micro stories used in preparing a martini.

```
If the problem is "Mix martini".
Intention: Mix martini.
The end.

If the intention is "Mix martini".
Step: Identify elements.
Step: Pour gin into glass.
Step: Pour vermouth into glass.
Step: Place lemon in glass.
The end.

If the intention is "Pour xx into yy".
Step: Put xx's container above yy.
Step: Tilt xx's container.
Step: Move xx's container to the table.
The end.
```

Note that we have added the notion of *step* to the problem-solving specifications suggested by Langley. Each step expresses a problem to be solved. The step sequences are micro stories that look much like recipes.

With the knowledge expressed this way, it is easy to have the problem solver ask for help when it gets stuck, to have a person supply suggestions, and even to learn how to perform various tasks by learning from the stories people tell when they go about making and fixing (Yang and Winston, In preparation).

4.3 Summary

One way to approach story understanding is to attack with all the reasoning, planning, representation, and architectural apparatus that Cognitive AI has to offer. Another approach is to presume that story understanding comes first, and that all the other Cognitive AI capabilities build on that story understanding.

5 Genesis illustrates, but obstacles block demonstration

We have prepared on the order of 1,000 stories for Genesis. Most are short, written for development, and just a few sentences in length, Others are much longer, some on the order of 100 sentences, and used for both development and demonstration.

The story genres include plays, novels, science fiction, stories from the Judeo-Christian Bible, law cases, social science and developmental psychology experiments, fairy tales, news accounts, Native American folktales, international conflict, and device repair.

When we add a new story to our story library, we sometimes find Genesis needs a new capability, but no new story has brought us to a stop.

Still, we consider all these to be illustrations, not demonstrations, because Genesis reads stories adapted by us from stories written by people for people. Genesis cannot read stories written by people for people without adaptation. That would constitute demonstration of a high order.

5.1 The internalization challenge

To get to demonstration of a high order, any story-understanding system must somehow internalize from the complex forms that writers employ and the ungrammatical, fragmentary English that people speak. With respect to this challenge, the recent work of Berwick and Chomsky has had a soothing effect. They say we have an external language because we have merge and we needed a way to communicate the resulting inner-language structures to others; they say we have so many languages because movement from and into our inner language is an engineering enterprise involving numerous arbitrary decisions (2016). From this perspective, our focus is and should be understanding what has to be in the inner language, and we can make progress as long as we can translate freely between that inner language and simple English. Understanding internalization/externalization at a human level is a different problem. We expect researchers focused on internalization/externalization will take on those challenges for us.

Lacking that internalization capability, some researchers work directly with stories written by people for people by treating the stories as bags of words to be processed by some sort of statistical method. Sometimes the statistical method involves deployment of a deep neural net. This sort of work can lead to powerful applications, but generally do not impress us as shedding much light on human story understanding.

Others, notably Mark Finlayson (2012), work with stories written by people for people by using a combination of programs that do synaptic analysis, word sense disambiguation, and the like, assisted by having human annotators help those programs over the difficult spots. We rewrite stories so as to not have difficult spots. Both approaches come with virtues and downsides.

We know that other big problems will likely emerge as progress is made on handling stories written by people for people, but for now, we think that is the major obstacle to convincing demonstration and practical application.

5.2 The knowledge acquisition challenge

Knowledge acquisition could emerge as one of those other big problems. Right now, we supply common-sense rules and concept patterns in English, as a parent or teacher would when explaining a story to a child or student. We have written on the order of 100 common-sense rules of various types—mostly deduction rules (∼70%) and enablement rules (∼20%)—along with a few dozen concept patterns, including, for example, *Revenge*, *Pyrrhic victory*, *Mistake because harmed*, *Mistake because unhappy*, *Suicide*, *Regicide*, *Answered prayer*, *Insane suicide*, *Retribution*, *Sell out*, *Aggression of a bully*, *Teaching a lesson*, and *Contradiction*.

Although we are writing what is needed in English, not programming it in, critics often complain that Genesis's knowledge should be learned, not provided.

We have mixed feelings about requiring everything to be learned. We think much of what we know comes from what we are told explicitly, not from learning grounded in what we have experienced, either directly or vicariously.

Nevertheless, self discovery has been and will remain prominent in our research plan. Mark Finlayson, for example, has done groundbreaking work on modeling Vladimir Propp's extraction, from Russian folk tales, of what we would call concept patterns (2012).

More recently, we have started to do research aimed at learning by being curious and asking questions. We imagine having human volunteers answering those questions via the Internet.

That kind of knowledge acquisition, acquisition by asking, would represent, we think, a middle ground between learning by being programmed and learning by processing millions of labeled examples.

Meanwhile, we have benefited by connecting Genesis to Lieberman's ConceptNet (Havasi *et al.*, 2009; Speer and Havasi, 2012), a connection that, in principle, increases the common-sense rules available to Genesis by two orders of magnitude.

Establishing a similar connection to Doug Lenat's Cyc system could also help solve the knowledge acquisition problem. We have not yet implemented that connection.

Still, such connections seem unlikely to satisfy all our common-sense needs because without appropriate prompting, human volunteers do not necessarily supply the kind of common sense that is so obvious, no one thinks to add it in the absence of the specific-case heuristic.

ConceptNet lists, for example, 43 causes of anger, including fishing, but does not offer an analog of a deduction rule that Genesis uses frequently, *If x harms y, then x angers y*.

5.3 The does-it-really-understand challenge

In a quick study of ConceptNet, Zhutian Yang explored eight verbs, extracting items interpretable as rules, with the following results, where rank identifies location in a list of most common English verbs (WordExample, 2018):

Verb	Rank	Rules
Leave	45	7
Win	93	35
Eat	111	98
Kill	183	101
Lie	191	59
Sleep	216	91
Drink	237	58
Cry	396	43

Taking a closer look at kill, Yang determined that of the approximately 100 rules obtained from ConceptNet, she could translate about half into plausibly useful Genesis rules. Many of what at first appeared to be deduction rules seemed better cast as explanation rules, leaving only the following five deduction rules:

```
If xx kills yy, then yy becomes dead.
If xx kills yy, then xx takes yy's life.
If xx kills yy, then xx ends yy's life.
If xx kills yy, then xx hurts yy's family.
If xx kills yy, then yy becomes pained.
```

The first of these deduction rules was already in Genesis. The second and third restate the first. The forth is much like *If x kills y and z is x's relative, then x harms z*, a rule also already in Genesis. And the fifth restates *If x kills y then x harms y*.

The remainder of the approximately 50 plausibly useful Genesis rules were explanation rules; the following are representative:

```
If xx wants to defend xx, then xx may kill yy.
If xx wants money, then xx may kill yy.
If xx is evil, then xx may kill yy.
If xx is mean, then xx may kill yy.
If xx wants to keep yy silent, then xx may kill yy.
```

Of the remaining items obtained from ConceptNet, many were redundant and some seemed of little value, such as these:

```
If xx wants to kill a cat, then xx may bury the cat.
If xx wants a nonviolent world, then xx may kill yy.
```

At the moment, the extraction of rules from ConceptNet requires manual editing. ConceptNet supplies a table with entries such as these:

Cause	Consequence
Defend yourself	Kill someone
Money	Kill
Be evil	Kill
Be mean	Kill someone
Keep them silent	Kill someone

Adding all the approximately 50 translated killing rules to Genesis produces no important change in what Genesis does with our test and demonstration stories; in the Crime Story, for example, Genesis adds four new elements for each killing via the additional deduction rules; none of the explanation rules was actuated.

Assuming the other verbs explored would produce similar deduction rule and explanation rule proportions, then on average each would produce about three or four deduction rules and approximately 30 rules all together. Doing this for, say, the most frequent 500 verbs, assuming each also produced about 30 rules, would yield 15,000 rules, two orders of magnitude more than we have provided to Genesis.

Would it matter much? Genesis would be able handle a richer assortment of stories, and if commanded, “Tell me about ‘killing’,” Genesis would start blurting out a lot of rules and seem to have a greater understanding of *killing*, but nothing new would be modeled.

5.4 The is-it-just-symbols challenge

No matter how many rules and concepts reside in Genesis, it remains reasonable to wonder if Genesis really understands killing or any other action. Lacking perceptual and motor apparatus, the symbols in any story-understanding system remain ungrounded in our physical world.

In general, we do not like to describe what we might implement, and everything we have described in section 3 is the product of an implemented system. Nevertheless, it seems fitting here to note that we are working on adding perceptual apparatus to Genesis so that a robot could tell the story of what it perceives and does as it wanders about, thus grounding such stories in the physical world.

5.5 Summary

Genesis's basic story understanding competences consist of common-sense rule deployment and concept search. These enable us to model many competences such as question answering, reading with controllable allegiances and cultural biases, personality trait modeling, trouble anticipation, conceptual similarity measurement, story alignment and analogical reasoning, question-driven interpretation, summarization, persuasive telling, new-story composition, and story-grounded hypothetical reasoning.

The development of all these implemented models constitute steps toward a computational account of human intelligence.

Nevertheless, all these are illustrations, not demonstrations, because we cannot yet deal with stories written by people for people.

6 Story understanding intersects with other fields

We suggest how progress in story-understanding research not only benefits from research in other fields, but also has much to offer.

6.1 Story understanding research suggests problems for Brain and Cognitive Science

Clearly there is much to be characterized. There are interesting opportunities to deploy experimental skills across the brain and cognitive spectrum. Questions must be sharpened, hypotheses formed, experiments devised, much data collected, results analyzed, mathematics developed—all leading to new hypotheses, new experiments, and perhaps new methods. Opportunities spring to mind for brain scanning, developmental studies, crowd-sourced validation of computational models, and probing supporting mechanisms, such as sequencing capabilities, in animal models. Here are a few representative examples of problems suggested by research in story understanding:

- Determine what small but essential elements we have in our brains that distinguish us from other species.
- Determine how symbolic descriptions can emerge from neural circuitry and architecture.
- Establish how and when our story-understanding competence develops in childhood.
- Establish how and when we develop common sense and conceptual knowledge in childhood and beyond.
- Establish how much we learn by being told versus by discovery.
- Establish what parts of our brains are the substrate for forming the sequences of classifications, properties, relations, actions, and events.
- Establish what parts of our brains are more or less engaged as story types vary.
- Understand how we internalize and process stories via vision and other senses, not just language, by studying what a visual story lights up in our brains that differs from what a linguistic story lights up.
- Understand whether programlike visual routines (Ullman, 1996) might arise from the same substrate that provides story understanding.
- Understand how we understand via empathetic thinking at every level from action recognition to reacting emotionally to the joys and griefs of others.

- Study everything else, because our story competence would be valueless without many competences we do, in fact, share with other species.

6.2 Story understanding research suggests challenges for research on neural nets

We are much impressed by the considerable success of deep neural nets in image classification. We also note that our heads are stuffed with neurons, and if you pluck them out, we do not think any more.

We resist, however, the supposition among some neural net enthusiasts that story understanding is best approached by working with hugely deep neural nets, equipped with billions of parameters, supplied with extremely large sample sets, mimicking what is done with image classification.

Such an approach puts the mechanism on top of the problem, an approach much criticized by David Marr, who argued against overenthusiastic attempts to use popular mechanisms on all problems. He believed that the problem to be understood should be on top and mechanisms should be selected only after a problem is understood computationally (1982).

A deep-net-on-top approach also seems at odds with the notion of abstraction barrier: once a system is worked out at one level, a higher level can use the resulting capabilities without reaching inside the lower level. Thinking in terms of abstraction suggests that we ask what capabilities must be implemented at the neural systems level to support the story understanding level.

We have suggested that basic story understanding requires only a small set of representational and computational capabilities. Representationally, there must be a way to construct complex, deeply nested symbolic descriptions; computationally, there must be ways to perform inference reflexes and do concept search. We believe those are the capabilities that neural-net researchers should try to model in biologically plausible neural systems.

6.3 Story understanding research has much to offer

Successful research in story understanding has much to offer because many fields contribute to an understanding of how stories are understood or to an understanding of how stories are put to use. The following are representative of the ways research in story understanding might have beneficial influence:

- For *linguistics*, suggestions about what must be in the inner language and what has to be externalized and internalized so as to enable a basic story-understanding competence.
- For the *humanities*, a characterization of what makes us quintessentially human.
- For *literary studies*, a way of thinking about how stories become structured, coherent, surprising, and memorable.
- For *design*, a starting point for thinking about how symbolic and visual thinking stimulate each other in creative action and retrospective critique.
- For *education*, new models of how people—in particular young children—think and learn, summarize and remember. These scientific advances may lead to practical advances in teaching, especially teaching through computational modeling.
- For fields in which previous experience is especially important, such as *politics*, *medicine*, *law*, *law enforcement*, *urban planning*, and *defense*, a promise of what-if tools that will assist and empower human experts in the same way spreadsheets assist and empower financial analysts.

- For the world, ways to mitigate some of the dangers that concern anxious futurists. By developing self-reflective architectures, story understanding offers machines that record their own experiences in story form, thereby enabling them to explain themselves—our only hope for safety as we come to depend more and more on our intelligent artifacts.

6.4 Summary

We believe story-understanding research has much to offer Brain and Cognitive Science because there is much to be learned about the neurobiology of story understanding and about how our ability to process stories develops epigenetically.

For research on neural nets, we suggest what capabilities should be modeled and that the models should be consistent with our growing understanding of neurobiology.

For other fields, with both scientific and engineering aims, we believe research on story understanding offers a way of thinking and the potential for new and exciting tools.

7 We benefit from inspirational and enabling research

In this section we identify representative research by others on story understanding, emphasizing research that has inspired and enabled our own.

7.1 We see much anticipation in Minsky's work

As we make progress, and go back to Marvin Minsky's work, we see much anticipation. Genesis exhibits, for example, aspects of all six of Minsky's levels of reasoning laid out in Part V of *Society of Mind* (2006):

- On the instinctive reaction and learned reaction levels, Genesis has common sense rules.
- On the deliberative thinking level, Genesis uses explanation rules where it sees opportunities to be more comprehending.
- On the reflective thinking level, Genesis finds concepts, such as *revenge*, in what it has done on lower levels.
- On the self-reflecting thinking level, work in progress will provide Genesis with a model of its own problem solving story.
- On the self-conscious reflection level, work in progress will provide Genesis with mental models of various actors in a story, what sorts of people they are, and what they know.

On the matter of self-conscious reflection, in Part IV of *Society of Mind* (1988), when discussing limits on what we can do simultaneously, Minsky wrote:

> ... we sometimes describe our thoughts as flowing in a 'stream of consciousness'—or as taking the form of an 'inner monologue' a process in which a sequence of thoughts seems to resemble a story or narrative.

We aspire to taking our understanding of that inner monologue to another level and likewise shed light on other aspects of the thinking Minsky wrote about in his seminal books (1988; 2006).

7.2 We salute the early pioneers

In this section we take note of research from which we drew inspiration and enablement, and we cite other research that focuses on related problems.

In section 2, we saluted the pioneering contributions of Schank, his students, and his colleagues (1972; 1977; 1981; 1991). Our work differs from theirs, in part, because we maintain libraries of previously encountered stories; Schank, in contrast, maintained a library of standardized descriptions, referred to as *scripts*, that depict the typical events involved, for example, in visiting a restaurant. Our work also differs from that of Schank in that Genesis translates relations and events into a rich inner language that consists largely of case frames in the style of Fillmore (1968), as described in section 2.2; Schank, in contrast, emphasized the translation of everything into a canonical language consisting of a dozen building blocks he called primitive acts.

In section 2.5, we discussed various sorts of common-sense rules. Of course there is nothing new in the use of deduction rules as that dates back at least to Aristotle. We noted that our post-hoc-ergo-propter-hoc rules are reminiscent of Charniak's demons (1972) and our censor rules were much discussed by Minsky (1988; 2006).

In section 2.9, we cited the work of Lehnert (1981). While one of Schank's students, she conceived the idea of *plot units*, which are somewhat like our concept patterns, except limited to combinations expressible in terms of mental states, positive events, and negative events.

7.3 We belong to a community of like-minded researchers

Though our story-enabled focus is unique, we are not alone. In this section, we highlight the work of some of our fellow travelers.

7.3.1 Script-based story competence

We are inspired by work in the field of computational narrative understanding and generation. Like Schank (Schank and Abelson, 1975; ?), we view story-processing as fundamental to human cognition, and we pursue language understanding not in isolation, but as part of a complete integrated cognitive system (Schank *et al.*, 1980b). Schank's frame-like script representations established a wide breadth of story-based applications, including personal and cultural memory, information retrieval (Schank *et al.*, 1980a), and interestingness (Schank, 1978). Throughout this script-based enterprise, we intersect with many of the early endeavors to model aspects of human story competence. Schank's (1973) MARGIE program filled commmonsense gaps, using stereotyped-script knowledge to answer comprehension questions. Through extensive use of scripts, Cullingford's (1978) SAM program paraphrased news stories, determined context through pattern-matching, and curtailed excessive inference (Cullingford, 1977), later even providing knowledge-based translations of its summaries (Carbonell *et al.*, 1981). Wilensky (1978; 1982) demonstrated the importance of narrative structure, goal-directedness, and failures with the plan-based PAM story-understanding program. Lehnert's QUALM program (Lehnert, 1977a,b) provided a cognitive theory of question-answering for story comprehension, augmenting both the SAM and PAM programs. Meehan's TALESPIN program (Meehan, 1976) generated realistic fable-like stories based on a complex interlocking model of personal drives, theory of mind, interpersonal relationships, and organization in space; though these representations and story-generation capabilities are possible within the Genesis framework, at present they remain outside of our direct focus.

Instead, our approach to story generation has mainly centered on alignment with previous stories (Fay, 2012) and judicious selection of sentences in an existing narrative to achieve a particular persuasive effect (Sayan, 2014). In our scaling-up efforts, we anticipate using Genesis's story-driven summarization capabilities to evaluate and compare the narrative quality of a large volumes of stories—produced either by Genesis or by an external story-generation system such as MEXICA (Pérez ý Pérez and Sharples, 2001). The news-reading programs FRUMP (DeJong and others, 1979; DeJong, 1982) and IPP (Lebowitz, 1981a,b) were some of the first to achieve significant breadth of coverage (processing and summarizing genuine news articles), though they did so by skimming at the expense of depth (deep understanding and rich cognitive or linguistic models). In our own scaling-up efforts, we aspire to a similar breadth of coverage and applications to real news stories. Carbonell's work on understanding politics and ideology in stories (Carbonell, 1978, 1979) informs aspects of our work ranging from culture-based mental models to the use of textual modulation as cues to authorial intent. Finally, the BORIS program (Lehnert *et al.*, 1981; Dyer, 1982; Dyer and Lehnert, 1982; Dyer, 1983) helped articulate the importance of understanding emotion and affect when making sense of stories, answering questions about them, and in particular deciding what to remember. In this regard, Lehnert's plot units (Lehnert, 1981) proved an effective representation for identifying abstract narrative themes and producing compressed stories, an approach strongly in line with Genesis's concept patterns.

In summary, though we do not adhere to script-based representations in particular, and though the skills possessed by these systems and the Genesis system do not wholly overlap, we are united by our approach to stories—broadly construed—as a fundamental substrate for cognition.

7.3.2 Planning, search, and inference

One of the key representations of the Genesis system is its use of various rule types to inject missing information or causal/inferential links, and concept patterns for identifying constellations of important thematic linkages in a story. In this regard, we are closely allied with formal reasoning and search subfields of AI: though Genesis does not use a formal reasoning engine to perform the necessary matching, firing, and search procedures, we nonetheless share an interest in supplying and efficiently collating common sense knowledge.

In our rule-based lineage, we recognize Hobbs as one of the early champions of this formal-logic approach to commonsense reasoning, e.g. through the TACITUS project (Hobbs *et al.*, 1987) for understanding stories about broken physical mechanisms through abductive reasoning. This formal-logic enterprise led to efforts to codify our everyday knowledge in formal systems and ontologies (Hobbs and Moore, 1985). These ontologies captured various aspects of life such as the qualitative behavior of liquids (Hayes, 1978), naive geography (Egenhofer and Mark, 1995), or spatial arrangement (Frank, 1997) and resonated with other non-computational cognitive models such as Gibson's (1979) uniquely naturalistic approach to vision. (The Genesis system's own rule-based commonsense knowledge is similarly expressed in the form of deductive and abductive rules, though in place of a uniform representation we have a heterogeneous variety of rule types, and we use other representations—such as transition-space diagrams—to capture other forms of knowledge. In this regard, we are closer to scruffy (free-form) approaches to knowledge engineering—less like CYC's (Lenat, 1995) formal expressions, more like OpenMind (Singh *et al.*, 2002) or ConceptNet (Liu and Singh, 2004; Speer *et al.*, 2008), knowledge bases which we have in fact linked into the Genesis system itself.) The formal-logic enterprise also spurred the development of techniques for escaping the apparent rigidity of deductive logic, including non-monotonic reasoning, particularly circumscription (McCarthy, 1981, 1986). In more recent efforts, etcetera abduction (Gordon, 2017) has provided a way to reify unstated commonsense assumptions as logical propositions and perform abductive or even probabilistic reasoning with them. By using etcetera abduction, Gordon has developed a model of commonsense human interactions which can infer goals, dispositions, and intents from a silent animation involving the motions of otherwise featureless geometric shapes (Gordon, 2016); the program emulates the human subjects of Heider and Simmel's (1944) famous psychological experiment. These variations and extensions to formal reasoning methods more closely match our computational-imperative-directed approach, according to which we do not adhere to a particular formalism but instead add whatever augmentations are necessary in order to capture the desired cognitive behavior.

Inference and formal logic of course place significant emphasis on the problem of search: knowing how to quickly uncover just the right inferences while precluding dead-ends. Though it bears repeating that the Genesis system does not currently make use of a formal planning system either in rule inference or problem solving, we draw inspiration from work attempting to capture large aspects of human cognition as forms of planning, search, and problem-solving. For example, Langley and others have produced an abduction-based account of diverse topics such as commonsense reasoning (Bridewell and Langley, 2011), problem-solving (Langley and Trivedi, 2013), interpersonal interaction (Meadows *et al.*, 2013), and personality-modulated conversation (Langley, 2017). Gordon, mentioned above, articulates the importance of representation to recognizing, storing, and retrieving plans (Gordon, 2004), while Riedl's programs (Riedl and Young, 2010) generate effective narratives through careful plot-driven planning. Many of these approaches regard planning as the fundamental substrate of cognition, where we regard planning as just another dimension of the story competence.

7.3.3 Architecture

We believe that if we are to model the full gamut of human intelligence, we must go beyond isolated expert systems, building tightly-integrated systems of interlocking modules which perform a variety of different tasks (Langley, 2006; Choi *et al.*, 2007). In this regard, we are driven by the directive articulated by Minsky (1988) to integrate a variety of methods, and by a well-established tradition of control-based systems whose functionality emerges by trading control between different subsystems rather than by purely top-down commands (Sussman, 1975; Longuet-Higgins, 1987; Sloman, 1993; Sloman and Poli, 1995; Larsen and Bundesen, 1996; Radul and Sussman, 2009). Representative architectures in the literature include Soar (Laird *et al.*, 1987; Laird, 2012), ACT-R (Anderson, 1996; Anderson *et al.*, 1997) with its focus on attention, CogPrime (Goertzel, 2012), ICARUS (Langley and Choi, 2006), EM-ONE (Singh and others, 2005) with its multiple layers of cognition, LIDA (Franklin and Patterson Jr, 2006; Snaider *et al.*, 2011) with its cognitive workspace approach, Polyscheme (Cassimatis, 2001), and Sigma (Rosenbloom, 2013). Our recent work on teaching Genesis problem-solving routines through natural language instruction is closely in line with the goals of the Rosie system, built atop Soar (Mohan *et al.*, 2012; Mohan and Laird, 2014; Mohan *et al.*, 2016), as well as other task-learning programs such as PLOW (Allen *et al.*, 2007).

7.3.4 Narratology and computational models of narrative

Because our research centers on the mechanisms underlying the human story competence, we are naturally interested in narratives themselves. We are enriched by the broad study of narratology and computational models of narrative. Fundamental narratological research—such as Propp's (Propp, 1968) systematic characterization of folk tales, Prince's (Prince, 1973, 1982) analysis of narrative-as-grammar, and Rimmon-Kena's (Rimmon-Kenan, 2003) account of what makes narratives different—all informs our understanding of narrative form.

Among the vast amount of literature on the computational front, there are a few areas of research that align especially closely with our own. The work of Vallas-Vargas and others (Valls-Vargas *et al.*, 2014) on identifying roles automatically in stories is of similar spirit to our own work (Finlayson, 2010, 2016) automatically inferring Propp-like categories from collections of folk tales (in contrast to using Propp's categories to generate new stories; see (Gervás, 2013)). Ryan (Ryan, 1991) and Michael (Michael, 2010) each articulate visions for computational approaches to narrative, and Michael's (Michael, 2012) formal-logic approach to capturing high-level narrative similarity resembles in focus, though not in implementation, our own model-based approach to narrative similarity (Winston, 2015). Charniak (Charniak, 1972) articulates the importance of commonsense knowledge to understanding stories (aspects of which we explore in more detail in the section on language), while Mueller provides a useful survey of the early story-understanding literature (Mueller, 2002) as well as several contributions to script-based (Mueller, 2004), formal-reasoning-based (Mueller, 2014), and multiple-representation-based (Mueller, 2003) approaches to commonsense reasoning in stories.

Our clearest point of departure from the narratology mainstream is that we are only secondarily concerned with narratives themselves; instead, we are primarily focused on studying how the mechanisms that enable us to understand stories likewise underly all uniquely human aspects of our intelligence.

7.3.5 Language

Though our use of the term *story competence* is highly general—encompassing our ability not only to comprehend written narratives, but engage in other processes such as symbolically representing complex visiospatial arrangements and integrating perception and language—our research in story understanding naturally and necessarily intersects with work on language. On natural language understanding, our approach is broadly similar to Winograd's (Winograd, 1972) in that Genesis's language-understanding apparatus is primarily directed toward gleaning meaning and only incidentally toward forming a parse tree. (See also, Lakoff and Thompson's points about the need for cognitive models of language (Lakoff and Thompson, 1975)). As with other existing systems, the Genesis system depends upon rich semantic knowledge of heterogeneous kinds, including WordNet taxonomies (Miller, 1995) and Conceptnet knowledge graphs (Liu and Singh, 2004). As we scale up toward a wider variety of language tasks and broader linguistic coverage, we expect the importance of such knowledge bases for deep semantic coverage will only increase. In our approach to engineering conversant programs—programs that can interact naturally with human users—we are fundamentally interested in building models of conversation that capture what humans do. In this sense, Genesis differs from chatterbots such as PARRY (Colby, 1981) or ELIZA (Weizenbaum, 1966) in which there is at most a shallow model of the cognitive processes underlying conversation. We are more in line with interactive systems such as PLOW (Allen *et al.*, 2007), which learns how to perform tasks in a web browser through human-computer dialogue. Question answering in particular constitutes its own sprawling subfield of human-computer interaction; Clark's knowledge-based approach (Clark *et al.*, 1999, 2016), which uses controlled English for breaking down questions and which empasizes knowledge-enabled flexibility and robustness, closely resembles our own. Though our work mostly eschews formal ontology, we are united by a shared interest in using a broad range of reasoning techniques to get decent performance on tasks such as, in Clark's case, science exams.

7.3.6 Special focus

Given our broad interest in studying general human cognition through the lens of story understanding, we have pursued particular projects that do not fit neatly within the other categories just surveyed.

We believe, for example, that in order to scale up to a large volume of efficiently-processable knowledge (an effort echoed by, for example, Gordon's work on learning from millions of personal stories (Gordon *et al.*, 2011)), we must be able to learn by analogy. In this regard, we are interested in the work of Barbella and colleagues on analogy in learning by reading (Barbella and Forbus, 2011; Friedman *et al.*, 2012), and on related work regarding conceptual change (Friedman and Forbus, 2010) and the role of rumination in the learning-by-reading process (Forbus *et al.*, 2007). Hofstadter (Hofstadter, 2001) and Lake et al. (Lake *et al.*, 2011) provide usefully different perspectives on analogy: Hofstadter positions analogy as the unifying principle of unique human intelligence, while Lake and colleagues implement a hierarchical probabilistic approach to learning the abstract unifying features of high-level concept categories.

Our work on hypothetical reasoning integrates a multiplicity of skills that have been studied before, such as predicting what characters want (Baker *et al.*, 2011), inferring character plans (Appelt and Pollack, 1992), and producing explanations under constraint (Magid *et al.*, 2015). In our work on hypothetical gap-filling (Holmes and Winston, 2016), we are informed by case-based expert reasoning Rissland (1989) as well as existing computational-cognitive models of hypothetical reasoning in physical domains Gerstenberg *et al.* (2017). Our hypothetical-based model of personality captures aspects of theory of mind. Like Baker et al. (2011), we model how an observer can accurately and effortlessly infer a character's beliefs and desires from the character's observed actions and use those inferences to make predictions. However, while Baker et al. cast the problem as inferring the joint distribution of a partially-observable Markov decision process that covers characters' beliefs, desires, and actions, we instead use a story-alignment process to infer implicit goals and constraints. In a similar way, Genesis's approach to hypothetical reasoning differs from other probabilistic approaches, e.g., Gerstenberg et al. (2017). The approaches are congruent in that they both emphasize the use of hypothetical alternatives as evidence. But whereas the model of Gerstenberg et al. makes use of an exact, step-by-step physics engine to predict what could probably happen, we uses qualitative behavioral models and story alignment to predict what can possibly happen. Following Sloman's (2017) observations, we posit that humans have knowledge and reasoning processes that are not easily captured by probability distributions over quantitatively-exact mechanical simulations, and that our partial, story-based surrogates of the world are a more plausible mechanism for predicting what characters can possibly do. Like Rizzo et al. (1997), we model variations in personality in terms of character goals and goal-directed actions. But where they emphasize goals as indicative of personality type, PERSONATE emphasizes constraints that force alternative *means* of achieving goals; the goals themselves are less characteristic of personality than the choice of means.

Because our personality-based program searches for the best unseen explanation for behavior, it can usefully be viewed as an abductive reasoner in the manner of (Meadows *et al.*, 2013). Similarly, the process of inferring characters' intentions and constraints can be viewed as a kind of plan ascription or recognition (Appelt and Pollack, 1992). Despite the simplicity of the one-step "plans" our program considers, the comparison with abductive reasoning is fruitful, but there are some differences in both approach and objective.

First, we model psychosocial causes and constraints, not general-purpose abductive inference. For this reason, our hypothetical reasoning module uses specialized procedures for developing, confirming, and rejecting behavioral hypotheses, rather than general-purpose inference processes. This approach is based on observations of how humans solve similar problems, with the special focus on explaining how it reaches its conclusions.

Second, unlike with more traditional abductive inference programs, the fundamental building blocks of Genesis's knowledge consist of compositional story-sequences, not atomic logical propositions (as assumed by the formal logic approach to counterfactuals taken by Ciardelli et al. (2016)). This design choice means that Genesis's knowledge has internal structure and that our module can reason using the full power of the underlying story-understanding system.

Because our module uses hypothetical reasoning to support commonsense inference in a story understanding context, our module is congruent with abductive approaches to commonsense reasoning such as Mueller's (2002). Importantly, however, our module is not limited to inferring events that occurred that were not explicitly mentioned (as in ordinary commonsense inference). Our module also considers counterfactual scenarios involving important events that were possible but never happened. These hypothetical alternatives play an important role in reasoning and take ordinary rule-based commonsense reasoning to another level.

7.3.7 Cognitive science

We are stimulated by work in the field of cognitive science, characterizing ourselves as approaching cognitive science from a computational modeling perspective. As noted earlier, our computational approach is entirely guided by our latest knowledge of what humans can do quickly, easily, and well; we are intrigued by systems that can replicate human fallacies and malfuctions. In this broad domain, we are especially inspired by the work of Spelke and others on core knowledge (Spelke and Kinzler, 2007) and the role of symbolic reasoning and language in visual problem solving (Hermer and Spelke, 1994). Chi (Chi *et al.*, 1981) demonstrates the importance of language processing in educational environment. Karmiloff-Smith (Karmiloff-Smith, 1995) lays the groundwork for the importance of conceptual change and redescription. Research by Schulz and others (Magid *et al.*, 2015; Gopnik and Schulz, 2004; Schulz and Gopnik, 2004; Schulz *et al.*, 2007) has provided the grist for our work on subjects as diverse as causal reasoning, conceptual change, and hypothetical reasoning.

7.3.8 Humanists

Being interested in human intelligence from all angles, we draw inspiration from fields beyond artificial intelligence and cognitive science. Our articulation of the inner language hypothesis is strongly aligned with paleoanthropologist Tattersall's work on the uniquely symbolic nature of the human species (Tattersall, 2012). Cobb's research (2013) on successful cultural mediation and Morris and Peng's work on cultural-dependent interpretation (1994) provided useful grounding for several of our projects centering on politics, worldview, negotiation, and diplomacy, and Goldstein's (1992) numerical scale for political conflict galvanized some of our own approaches to quantifying harm and escalation. Less directly, diverse work—such as Austin's (Austin, 1956) on how excuses explain moral accidents, Boden's on creativity (Boden, 2009), Mueller's on daydreaming (Mueller, 1990), Dennet's on theory of mind (Daniel, 1991), Lakoff and Johnson's on the primacy of metaphor in cognition (Lakoff and Johnson, 2008), Hurley and colleagues' on humor (Hurley *et al.*, 2011), and Charon's (2008) on applying narrative theory to providing better continuity of medical care—reaffirms our belief in the centrality of narrative to the human condition.

7.4 Summary

Prior research has inspired and enabled our own. We are inspired, for example, by Marvin Minsky's work on levels of reasoning as well as enabled by prior research in story understanding by Schank and many others.

8 Contributions

We have explained why we are passionate about developing computational models of our human ability to understand stories. In particular, we have:

- Noted that the ability to assemble complex, highly nested symbolic descriptions enabled our human story competences. Like the keystone in an arch, that symbolic ability derives value from and makes more valuable other capabilities already present, such as those involved in perception and sequence remembering.
- Explained that our goal is to build the simplest possible system that exhibits humanlike behaviors, not to build an applied question-answering system or to advance the state of the art in computational or representational sophistication. Because we do not want a system that exhibits unnatural behaviors, we adhere to the *computational imperative principle*, including in Genesis only what it needs to exhibit humanlike behaviors.
- Exhibited many kinds of common-sense rules and concept patterns, showing how they are put to work in *inference reflexes* and *concept recognition* so as to support models of many story-understanding competences, enabling Genesis to answer questions, describe conceptual content, summarize, compare and contrast stories, react with cultural biases, instruct, reason hypothetically, solve problems, and find useful precedents.
- Suggested that our human problem-solving abilities are enabled by our story understanding competences, rather than vice versa.
- Discussed how research in story understanding benefits from and can benefit many fields of science and application.

We do what we do because of out-of-control scientific curiosity, of course, but we anticipate that the scientific answers will revolutionize the engineering of intelligent systems. If we develop a top-to-bottom account of our story-understanding competences, applications with human-like intelligence and self awareness will emerge and empower in education, economics, politics, health care, law, law enforcement, urban planning, defense, and business. Some of those applications will be linked together in analogs of social networks, opening up world-changing opportunities in energy, the environment, cybersecurity, and other high-impact areas with otherwise unsolvable problems.

9 Acknowledgements

The work described in this paper was supported, in part, by the Air Force Office of Scientific Research, contract number FA9550-17-1-0081. Other sponsors of work on aspects of the Genesis Group's work include the Office of Naval Research, award number N00014091059, and the National Science Foundation, via the Center for Brains, Minds, and Machines, STC award CCF-1231216.

We thank Robert C. Berwick, Christopher Fry, Gabriel Kreiman, Joe Lewis and Nick Montfort for important suggestions. We note that we have not addressed all of their suggestions and that making a suggestion is not an endorsement of all the thoughts we express.

References

James Allen, Nathanael Chambers, George Ferguson, Lucian Galescu, Hyuckchul Jung, Mary Swift, and William Taysom. PLOW: A collaborative task learning agent. In *AAAI*, volume 7, pages 1514–1519, 2007.

John R. Anderson, Michael Matessa, and Christian Lebiere. Act-r: A theory of higher level cognition and its relation to visual attention. *Human-Computer Interaction*, 12(4):439–462, 1997.

John R. Anderson. ACT: A simple theory of complex cognition. *American Psychologist*, 51(4):355, 1996.

Douglas E. Appelt and Martha E. Pollack. Weighted abduction for plan ascription. *User modeling and user-adapted interaction*, 2(1):1–25, 1992.

John L. Austin. A plea for excuses: The presidential address. In *Proceedings of the aristotelian society*, volume 57, pages 1–30. JSTOR, 1956.

Hiba Awad. Culturally based story understanding. Master's thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 2013.

Yu Bai, Chris Pearce, Pat Langley, Mike Barley, and Charlotte Worsfold. An architecture for flexibly interleaving planning and execution. *Advances in Cognitive Systems*, 2015.

Chris Baker, Rebecca Saxe, and Joshua Tenenbaum. Bayesian theory of mind: Modeling joint belief-desire attribution. In *Proceedings of the Cognitive Science Society*, volume 33, 2011.

David Michael Barbella and Kenneth D. Forbus. Analogical dialogue acts: Supporting learning by reading analogies in instructional texts. In *AAAI*, 2011.

Robert C. Berwick and Noam Chomsky. *Why Only Us*. MIT Press, Cambridge, MA, 2016.

Margaret A. Boden. Computer models of creativity. *AI Magazine*, 30(3):23, 2009.

Will Bridewell and Pat Langley. A computational account of everyday abductive inference. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 33, 2011.

Jaime G. Carbonell, Richard Edward Cullingford, and Anatole V. Gershman. Steps toward knowledge-based machine translation. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (4):376–392, 1981.

Jaime G. Carbonell. Politics: Automated ideological reasoning. *Cognitive Science*, 2(1):27–51, 1978.

Jaime G. Carbonell. Subjective understanding: Computer models of belief systems. Technical report, Department of Computer Science, Yale University, 1979.

Nicholas Louis Cassimatis. *Polyscheme: a cognitive architecture for intergrating multiple representation and inference schemes*. PhD thesis, Massachusetts Institute of Technology, 2001.

Eugene Charniak. *Toward a model of children⊠s story comprehension*. PhD thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 1972.

Rita Charon. *Narrative medicine: Honoring the stories of illness*. Oxford University Press, 2008.

Michelene Chi, Paul J. Feltovich, and Robert Glaser. Categorization and representation of physics problems by experts and novices. *Cognitive science*, 5(2):121–152, 1981.

Dongkyu Choi, Tolgo Konik, Negin Nejati, Chunki Park, and Pat Langley. Structural transfer of cognitive skills. In *Proceedings of the eighth international conference on cognitive modeling*, pages 115–120. Citeseer, 2007.

Ivano Ciardelli, Linmin Zhang, and Lucas Champollion. Two switches in the theory of counterfactuals: A study of truth conditionality and minimal change. 2016.

Peter Clark, John Thompson, and Bruce Porter. A knowledge-based approach to question-answering. In *Proc. AAAI*, volume 99, pages 43–51. Citeseer, 1999.

Peter Clark, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter D. Turney, and Daniel Khashabi. Combining retrieval, statistics, and inference to answer elementary science questions. In *AAAI*, pages 2580–2586, 2016.

Sara B. Cobb. *Speaking of violence: The politics and poetics of narrative in conflict resolution*.

Oxford University Press, 2013.

Kenneth Mark Colby. Modeling a paranoid mind. *Behavioral and Brain Sciences*, 4(4):515–534, 1981.

Richard Edward Cullingford. Controlling inference in story understanding. In *IJCAI*, volume 77, page 17, 1977.

Richard Edward Cullingford. Script application: computer understanding of newspaper stories. Technical report, Department of Computer Science, Yale University, 1978.

Dennett Daniel. *Consciousness explained*. New York: Little Brown & Co, 1991.

Gerald DeJong et al. Skimming stories in real time: An experiment in integrated understanding. Technical report, Department of Computer Science, Yale University, 1979.

Gerald DeJong. An overview of the FRUMP system. *Strategies for natural language processing*, 113:149–176, 1982.

Michael G. Dyer and Wendy G. Lehnert. Question answering for narrative memory. In *Advances in Psychology*, volume 9, pages 339–358. Elsevier, 1982.

Michael G. Dyer. In-depth understanding. a computer model of integrated processing for narrative comprehension. Technical report, Department of Computer Science, Yale University, 1982.

Michael G. Dyer. The role of affect in narratives. *Cognitive Science*, 7(3):211–242, 1983.

Max J. Egenhofer and David M. Mark. Naive geography. In *International Conference on Spatial Information Theory*, pages 1–15. Springer, 1995.

Matthew Paul Fay. Enabling imagination through story alignment. Master's thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 2012.

Matthew Paul Fay. *Story Understanding in Genesis: Exploring Automatic Plot Construction through Commonsense Reasoning*. PhD thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 2014.

Evelina Fedorenko and Rosmary Varley. Language and thought are not the same thing: evidence from neuroimaging and neurological patients. *Annals of the New York Academy of Sciences*, 2016.

Christiane Fellbaum, editor. *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge, MA, 1998.

David Ferrucci, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya Kalyanpur, Adam Lally, J. William Murdock, Eric Nyberg, John Prager, Nico Schlaefer, and Chris Welty. The AI behind watson—the technical article. *AI Magazine*, 2010.

Charles J. Fillmore. The case for case. In Emmon Bach and R. Harms, editors, *Universals in Linguistic Theory*, pages 1–88. Holt, Rinehart, and Winston, 1968.

Mark Alan Finlayson. Learning narrative morphologies from annotated folktales. In *Proceedings of the 1st International Workshop on Automated Motif Discovery in Cultural Heritage and Scientific Communication Texts*, 2010.

Mark Alan Finlayson. *Learning Narrative Structure from Annotated Folktales*. PhD thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 2012.

Mark Alan Finlayson. Inferring propp's functions from semantically annotated text. *The Journal of American Folklore*, 129(511):55–77, 2016.

Kenneth D. Forbus, Christopher Riesbeck, Lawrence Birnbaum, Kevin Livingston, Abhishek Sharma, and Leo Ureel. Integrating natural language, knowledge representation and reasoning, and analogical processing to learn by reading. In *Proceedings of the National Conference on Artificial Intelligence*, volume 22, page 1542. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2007.

Andrew U. Frank. Spatial ontology: A geographical information point of view. In *Spatial and temporal reasoning*, pages 135–153. Springer, 1997.

Stan Franklin and F.G. Patterson Jr. The LIDA architecture: Adding new modes of learning to an intelligent, autonomous, software agent. *pat*, 703:764–1004, 2006.

Scott E. Friedman and Kenneth D. Forbus. An integrated systems approach to explanation-based conceptual change. In *AAAI*, 2010.

Scott E. Friedman, David Michael Barbella, and Kenneth D. Forbus. Revising domain knowledge with cross-domain analogy. *Advances in Cognitive Systems*, 2:13–24, 2012.

Tobias Gerstenberg, Liang Zhou, Kevin A. Smith, and Joshua B. Tenenbaum. Faulty towers: A hypothetical simulation model of physical support. In Glenn Gunzelmann, Andrew Howes, Thora Tenbrink, and Eddy Davelaar, editors, *Proceedings of the 39th Annual Conference of the Cognitive Science Society*, pages 409–414, Austin, TX, 2017. Cognitive Science Society.

Pablo Gervás. Propp's morphology of the folk tale as a grammar for generation. In *OASIcs-OpenAccess Series in Informatics*, volume 32. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2013.

James J. Gibson. *The Ecological Approach to Visual Perception*. 1979.

Clark Glymour. Bayesian ptolemaic psychology. In William Harper and Gregory Wheeler, editors, *Probability and Inference: Essays in Honour of Henry E. Kyburg, Jr.* King's College Publilcations, London, 2007.

Ben Goertzel. Cogprime: An integrative architecture for embodied artificial general intelligence. *Dynamical Psychology: An International, Interdisciplinary Journal of Complex Mental Processes*, 2012.

Joshua S. Goldstein. A conflict-cooperation scale for weis events data. *Journal of Conflict Resolution*, 36(2), June 1992.

Alison Gopnik and Laura E. Schulz. Mechanisms of theory formation in young children. *Trends in cognitive sciences*, 8(8):371–377, 2004.

Andrew S. Gordon, Cosmin Adrian Bejan, and Kenji Sagae. Commonsense causal reasoning using millions of personal stories. In *AAAI*, 2011.

Andrew S. Gordon. The representation of planning strategies. *Artificial Intelligence*, 153(1-2):287–305, 2004.

Andrew S. Gordon. Commonsense interpretation of triangle behavior. In *AAAI*, pages 3719–3725, 2016.

Andrew S. Gordon. Solving interpretation problems with etcetera abduction. In *Proceedings of the Fifth Annual Conference on Advances in Cognitive Systems*, Troy, NY, 2017.

Jonathan Gottschall. *The Storytelling Animal: How Stories Make Us Human*. Houghton Mifflin Harcourt, Boston, MA, 2012.

H. Paul Grice. *Studies in the Way of Words*. Harvard University Press, Cambridge, MA, 1989.

Catherine Havasi, James Pustejovsky, Robert Speer, and Henry Lieberman. Digital intuition: applying common sense using dimensionality reduction. *IEEE Intelligent Systems*, 24(4):24–35, 2009.

Patrick J. Hayes. The naive physics manifesto. *Expert systems in the micro-electronic age*, 1978.

Fritz Heider and Marianne Simmel. An experimental study of apparent behavior. *The American journal of psychology*, 57(2):243–259, 1944.

Linda Hermer and Elizabeth S. Spelke. A geometric process for spatial reorientation in young children. *Nature*, 370(6484):57, 1994.

Jerry R. Hobbs and Robert C. Moore. *Formal theories of the commonsense world*, volume 1. Intellect Books, 1985.

Jerry R. Hobbs, William Croft, Todd Davies, Douglas Edwards, and Kenneth Laws. Commonsense metaphysics and lexical semantics. *Computational linguistics*, 13(3-4):241–250, 1987.

Douglas R. Hofstadter. Analogy as the core of cognition. *The analogical mind: Perspectives from cognitive science*, pages 499–538, 2001.

Dylan Holmes and Patrick Winston. Story-enabled hypothetical reasoning. In *Proceedings of the Fourth Annual Conference on Advances in Cognitive Systems*, June 2016.

Matthew M. Hurley, Daniel Clement Dennett, and Reginald B. Adams Jr. *Inside jokes: Using humor to reverse-engineer the mind*. MIT press, 2011.

A. Karmiloff-Smith. *Beyond Modularity: A Developmental Perspective on Cognitive Science*. A Bradford Book, 1995.

Boris Katz. Annotating the World Wide Web using natural language. In *Proceedings of the 5th RIAO Conference on Computer Assisted Information Searching on the Internet*, pages 136–159, 1997.

Caryn Krakauer. Story retrieval and comparison using concept patterns. Master's thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 2012.

John E. Laird, Allen Newell, and Paul S. Rosenbloom. Soar: An architecture for general intelligence. *Artificial intelligence*, 33(1):1–64, 1987.

John E. Laird. *The Soar cognitive architecture*. MIT press, 2012.

Brenden Lake, Ruslan Salakhutdinov, Jason Gross, and Joshua Tenenbaum. One shot learning of simple visual concepts. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 33, 2011.

George Lakoff and Mark Johnson. *Metaphors we live by*. University of Chicago press, 2008.

George Lakoff and Henry Thompson. Introducing cognitive grammar. In *Annual Meeting of the Berkeley Linguistics Society*, volume 1, pages 295–313, 1975.

Pat Langley and Dongkyu Choi. A unified cognitive architecture for physical agents. In *Proceedings of the National Conference on Artificial Intelligence*, volume 21, page 1469. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2006.

Pat Langley and Nishant Trivedi. Elaborations on a theory of human problem solving. *Advances in Cognitive Systems*, 3:1–12, 2013.

Pat Langley, Miranda Emery, Michael Barley, and Christopher J. MacLellan. An architecture for flexible problem solving. *Advances in Cognitive Systems*, 2013.

Pat Langley, Chris Pearce, Mike Barley, and Miranda Emery. Bounded rationality in problem solving: Guiding search with domain-independent heuristics. *Mind and Society*, 13, 2014.

Pat Langley, Chris Pearce, Yu Bai, Charlotte Worsfold, and Mike Barley. Variations on a theory of problem solving. *Advances in Cognitive Systems*, 2016.

Pat Langley. Cognitive architectures and general intelligent systems. *AI magazine*, 27(2):33, 2006.

Pat Langley. A cognitive systems analysis of personality and conversational style. In *Proceedings of the Fifth Annual Conference on Advances in Cognitive Systems*, June 2017.

Axel Larsen and Claus Bundesen. A template-matching pandemonium recognizes unconstrained handwritten characters with high accuracy. *Memory & cognition*, 24(2):136–143, 1996.

Michael Lebowitz. Generalization and memory in an integrated understanding system. 1981.

Michael Lebowitz. The nature of generalization in understanding. 1981.

Wendy G. Lehnert, Michael G. Dyer, Peter N. Johnson, C.J. Yang, and Steve Harley. Boris–an experiment in in-depth understanding of narratives. Technical report, Department of Computer Science, Yale University, 1981.

Wendy G. Lehnert. A conceptual theory of question answering. In *Proceedings of the 5th international joint conference on Artificial intelligence-Volume 1*, pages 158–164. Morgan Kaufmann Publishers Inc., 1977.

Wendy G. Lehnert. The process of question answering. Technical report, Department of Computer Science, Yale University, 1977.

Wendy G. Lehnert. Plot units and narrative summarization. *Cognitive Science*, 5(4):293–331, 1981.

Douglas B. Lenat. Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38, 1995.

Hugo Liu and Pushpinder Singh. Conceptnet—a practical commonsense reasoning tool-kit. *BT technology journal*, 22(4):211–226, 2004.

H. Christopher Longuet-Higgins. *Mental processes: Studies in cognitive science*. The MIT Press, 1987.

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

David Marr. Artificial intelligence—a personal view. *Artificial Intelligence*, 9:37–48, 1977.

David Marr. *Vision*. W. H. Freeman, San Francisco, CA, 1982.

John McCarthy. Circumscription—a form of non-monotonic reasoning. In *Readings in Artificial Intelligence*, pages 466–472. Elsevier, 1981.

John McCarthy. Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence*, 28:89–116, 1986.

Ben Meadows, Pat Langley, and Miranda Emery. Understanding social interactions using incremental abductive inference. In *Proceedings of the Second Annual Conference on Advances in Cognitive Systems ACS*, volume 39, page 56, 2013.

James Richard Meehan. The metanovel: writing stories by computer. Technical report, Department of Computer Science, Yale University, 1976.

Loizos Michael. Computability of narrative. In *AAAI Fall Symposium: Computational Models of Narrative*, 2010.

Loizos Michael. Similarity of narratives. In *Proceedings of the LREC 2012 Workshop on Computational Models of Narrative*, 2012.

George A. Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995.

Marvin Minsky. *The Society of Mind*. Simon and Schuster, New York, NY, 1988.

Marvin Minsky. *The Emotion Machine*. 2006.

Shiwali Mohan and John E. Laird. Learning goal-oriented hierarchical tasks from situated interactive instruction. In *AAAI*, pages 387–394, 2014.

Shiwali Mohan, Aaron H. Mininger, James R. Kirk, and John E. Laird. Acquiring grounded representations of words with situated interactive instruction. In *Advances in Cognitive Systems*. Citeseer, 2012.

Shiwali Mohan, James Kirk, and John E. Laird. A computational model for situated task learning with interactive instruction. *arXiv preprint arXiv:1604.06849*, 2016.

Michael W. Morris and Kaiping Peng. Culture and cause: American and Chinese attributions for social and physical events. *Journal of Personality and Social Psychology*, 67(6):949–971, 1994.

Erik T. Mueller. *Daydreaming in humans and machines: a computer model of the stream of thought*. Intellect Books, 1990.

Erik T. Mueller. Story understanding. *Encyclopedia of Cognitive Science*, 2002.

Erik T. Mueller. Story understanding through multi-representation model construction. In *Proceedings of the HLT-NAACL 2003 workshop on Text meaning-Volume 9*, pages 46–53. Association for Computational Linguistics, 2003.

Erik T. Mueller. Understanding script-based stories using commonsense reasoning. *Cognitive Systems Research*, 5(4):307–340, 2004.

Erik T. Mueller. *Commonsense reasoning: an event calculus based approach*. Morgan Kaufmann, 2014.

David Nackoul. Text to text: Plot unit searches generated from English. Master's thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 2010.

Jessica Noss. Who knows what? perspective-enabled story understanding. Master's thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 2017.

Rafael Pérez ý Pérez and Mike Sharples. MEXICA: A computer model of a cognitive account of creative writing. *Journal of Experimental & Theoretical Artificial Intelligence*, 13(2):119–139, 2001.

Livia Polanyi. *Telling the American Story: A structural and Cultural Analysis of Conversational Storytelling*. MIT Press, Cambridge, MA, 1989.

Gerald Prince. A grammar of stories. *The Hague: Mouton*, 1973.

Gerald Prince. *Narratology: The Form and Functioning of Narrative*. Janua Linguarum : studia memoriae Nicolai van Wijk dedicata. Series maior. Mouton, 1982.

Vladimir Propp. *Morphology of the folktale*. University of Texas Press, 1968.

Alexey Radul and Gerald Jay Sussman. The art of the propagator. In *Proceedings of the 2009 international lisp conference*, pages 1–10, 2009.

Mark O. Riedl and Robert Michael Young. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research*, 39:217–268, 2010.

Shlomith Rimmon-Kenan. *Narrative fiction: Contemporary poetics*. London:Routledge, 2003.

E. L. Rissland. Dimension-based analysis of hypotheticals from supreme court oral argument. In *Proceedings of the 2nd International Conference on Artificial Intelligence and Law*, ICAIL '89, pages 111–120, New York, NY, USA, 1989. ACM.

Paola Rizzo, Manuela Veloso, Maria Miceli, and Amedeo Cesta. Personality-driven social behaviors in believable agents. In *Proceedings of the AAAI Fall Symposium on Socially Intelligent Agents*, pages 109–114, 1997.

Paul S. Rosenbloom. The sigma cognitive architecture and system. *AISB Quarterly*, 136:4–13, 2013.

Marie-Laure Ryan. *Possible worlds, artificial intelligence, and narrative theory*. Indiana University Press, 1991.

Sila Sayan. Audience aware computational discourse generation for instruction and persuasion. Master's thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 2014.

Roger C. Schank and Robert P. Abelson. Scripts, plans, and knowledge. In *IJCAI*, pages 151–157, 1975.

Roger C. Schank and Robert P. Abelson. *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures*. Lawrence Erlbaum Associates, Mahwah, NJ, 1977.

Roger C. Schank and Christopher K. Riesbeck. *Inside Computer Understanding: Five Programs Plus Miniatures*. Lawrence Erlbaum Associates, Mahwah, NJ, 1981.

Roger C. Schank, Neil M. Goldman, Charles J. Rieger III, and Christopher Riesbeck. Margie: Memory analysis response generation, and inference on english. In *IJCAI*, volume 3, pages 255–261, 1973.

Roger C. Schank, Janet L. Kolodner, and Gerald DeJong. Conceptual information retrieval. In *Proceedings of the 3rd annual ACM conference on Research and development in information retrieval*, pages 94–116. Butterworth & Co., 1980.

Roger C. Schank, Michael Lebowitz, and Lawrence Birnbaum. An integrated understander. *Computational Linguistics*, 6(1):13–30, 1980.

Roger C. Schank. Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology*, 3(4):552–631, 1972.

Roger C. Schank. Interestingness: Controlling inferences. Technical report, Department of Computer Science, Yale University, 1978.

Roger C. Schank. *Tell Me a Story: A New Look at Real and Artificial Intelligence*. Northwestern University Press, Evanston, IL, 1991.

Laura E. Schulz and Alison Gopnik. Causal learning across domains. *Developmental psychology*, 40(2):162, 2004.

Laura E. 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, 2007.

Pushpinder Singh et al. *EM-ONE: an architecture for reflective commonsense thinking*. PhD thesis, Massachusetts Institute of Technology, 2005.

Pushpinder Singh, Thomas Lin, Erik T. Mueller, Grace Lim, Travell Perkins, and Wan Li Zhu. Open mind common sense: Knowledge acquisition from the general public. In *OTM Confederated International Conferences: On the Move to Meaningful Internet Systems"*, pages 1223–1237. Springer, 2002.

Aaron Sloman and Riccardo Poli. Sim_agent: A toolkit for exploring agent designs. In *International Workshop on Agent Theories, Architectures, and Languages*, pages 392–407. Springer, 1995.

Aaron Sloman. The mind as a control system. *Royal Institute of Philosophy Supplements*, 34:69–110, 1993.

Aaron Sloman. Huge but unnoticed gaps between current AI and natural intelligence. In *Philosophy and Theory of Artificial Intelligence Conference*, November 2017.

Javier Snaider, Ryan McCall, and Stan Franklin. The LIDA framework as a general tool for AGI. In *International Conference on Artificial General Intelligence*, pages 133–142. Springer, 2011.

Susan S. Song. Of intent and action: Implementing personality traits for storytelling through concept patterns. Master's thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 2012.

Robert Speer and Catherine Havasi. Representing general relational knowledge in conceptnet 5. In *Proceedings of LREC*, 2012. Available via http://lrec-conf.org/proceedings/lrec2012/pdf/1072_Paper.pdf.

Robert Speer, Catherine Havasi, and Henry Lieberman. Analogyspace: Reducing the dimensionality of common sense knowledge. In *AAAI*, volume 8, pages 548–553, 2008.

Elizabeth S. Spelke and Katherine D. Kinzler. Core knowledge. *Developmental science*, 10(1):89–96, 2007.

Gerald Jay Sussman. *A computer model of skill acquisition*. Elsevier Science Inc., 1975.

Ian Tattersall. *Becoming Human*. Harcourt, Orlando, FL, 1998.

Ian Tattersall. Human evolution and cognition. *Theory in Biosciences*, 123(4):371–379, 2010.

Ian Tattersall. *Masters of the Planet: the search for our human origins*. Palgrave, Macmillan, 2012.

Shimon Ullman. *High-Level Vision*. MIT Press, Cambridge, MA, 1996.

Josep Valls-Vargas, Jichen Zhu, and Santiago Ontanón. Toward automatic role identification in unannotated folk tales. In *Tenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2014.

Joseph Weizenbaum. ELIZA–a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1):36–45, 1966.

Robert Wilensky. Understanding goal-based stories. Technical report, Department of Computer Science, Yale University, 1978.

Robert Wilensky. Points: A theory of the structure of stories in memory. *Strategies for natural language processing*, 345:374, 1982.

Terry Winograd. Understanding natural language. *Cognitive psychology*, 3(1):1–191, 1972.

Patrick Henry Winston. *Artificial Intelligence*. Addison-Wesley, 3 edition, 1992.

Patrick Henry Winston. The strong story hypothesis and the directed perception hypothesis. *AAAI*, 2011.

Patrick Henry Winston. The next 50 years: a personal view. *Biologically Inspired Cognitive Architectures*, 2012.

Patrick Henry Winston. The right way. *Cognitive Systems Foundation*, 2012.

Patrick Henry Winston. Model-based story summary. In Mark Alan Finlayson, Ben Miller, and Remi Ronfard, editors, *Proceedings of the 6th Workshop on Computational Models of Narrative(CMN 2015)*, volume 45. Oasics, 2015.

Patrick Henry Winston. Self-aware problem solving. In preparation.

WordExample. Most common verbs in english (1,000), 2018. https://www.wordexample.com/list/most-common-verbs-english.

Zhutian Yang and Patrick Henry Winston. Learning by asking questions and learning by aligning stories: how a story-grounded problem solver can acquire knowledge. In preparation.

Charles Yang. Ontogeny and phylogeny of language. *Proceedings of the National Academy of Sciences*, 110(16):177–208, 2013.

Wolfgang Victor Hayden Yarlott. Old man coyote stories: Cross-cultural story understanding in the genesis story understanding system. Master's thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 2014.