

The Genesis Manifesto: Story Understanding and Human Intelligence

Patrick Henry Winston and Dylan Holmes

May 4, 2017

This is β draft #4. We plan to supply additional references and to address a long list of questions and suggestions. We also plan to describe new work completed after January, 2017

Abstract

We believe we must construct biologically plausible computational models of human story understanding if we are to develop a computational account of human intelligence. We argue that building a story-understanding system exposes computational imperatives associated with human competences such as question answering, mental modeling, culturally biased story interpretation, story-based hypothetical reasoning, and self-aware problem solving. We explain that we believe such human competences rest on a uniquely human ability to construct complex, highly nested symbolic descriptions.

We illustrate our approach to modeling human story understanding by describing the development of the Genesis story understanding system and by explaining how Genesis goes about understanding short, 20- to 100-sentence stories expressed in English. The stories include, for example, summaries of plays, such as Shakespeare's *Macbeth*; fairy tales, such as *Hansel and Gretel*; and contemporary conflicts, such as the 2007 Estonia–Russia cyberwar.

We explain how we ensure that work on Genesis is scientifically grounded, we identify representative questions to be answered by empirical science, and we note why story understanding has much to offer not only to Artificial Intelligence but also to fields such as business, defense, design, economics, education, humanities, law, linguistics, neuroscience, philosophy, psychology, medicine, and politics.

Keywords: computational models of human intelligence; story understanding; merge-enabled description; computational imperatives; inference reflexes; concept discovery; Genesis story-understanding system.

1 Vision

Our goal is to develop a comprehensive computational account of human intelligence. To develop such an account, 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 are the symbolic species and that being symbolic also enabled us to become the story understanding species. Our answer to the support-and-benefit question is that our symbolic competence, and the story competence that it enables, could not have evolved without myriad elements already in place.

Our purpose in the rest of this section is to explain what being symbolic means, how being symbolic enables story understanding, what exactly we mean when we write that we are studying story understanding, how we approach our study of story understanding, and why our approach is scientifically grounded. Our work on the support-and-benefit question constitutes another story that is not yet ready to be told.

We begin by noting recent claims by Robert Berwick and Noam Chomsky. In *Why Only Us* (2016), they argue that only we 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 just one small step for evolution, we hypothesize that it enables a giant leap in intelligence, because merge gives us the ability to build complex, highly nested symbolic descriptions of classes, properties, relations, actions, and events.

Having that ability to build complex, highly nested symbolic descriptions is what being symbolic means. With that ability 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.

We claim that our 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. We give a name to collections of these 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, for example, causal, means-ends, 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, but when we refer to an *inner story*, we refer to a story expressed in an inner, merge-enabled language that rats and dogs and chimpanzees either do not have or do not have on anything like our level. Note also that we include what some narratologists would exclude, because 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).

1.1 The Strong Story Hypothesis

We have argued that we humans have the ability to build complex, highly nested symbolic descriptions, perhaps via a Berwick–Chomsky *merge* operation, perhaps via some other triumph of evolution. That symbolic ability enabled the composition and exploitation of inner stories. Eventually, we developed the means to externalize those inner stories and to internalize stories presented to us, and because we are social animals, externalization and internalization had a powerful amplifying effect. Because all these abilities seem unique to our species, at least at anything like our level, Winston introduced 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).

We think telling, understanding, and recombining have immense consequences. We think it reasonable, for example, to view recipe following as a special case of story understanding. Reasoning is a special case of recipe following. Planning is a special case of reasoning. Education begins with listening to fairy tales, many of which aim to scare us into behaving properly; then, we acquire precedent stories for later use packaged up in subjects such as history, literature, law, medicine, business, engineering, science, and religion; learning how to think is a matter of acquiring the skills involved in deploying such precedent stories. Creativity relies heavily on finding and recombining fragments drawn from precedents.

Note that we have no desire to be human chauvinists. Of course, 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 bees to chimpanzees exhibit impressive capabilities, and we agree that those capabilities are evidence of various kinds of intelligence.

1.2 We start by specifying the behavior we want to understand

Our approach to developing models of story understanding is highly influenced by Marvin Minsky's notion of suitcase word; that is, words such as *intelligence*, *creativity*, *consciousness*, and *story understanding* 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). Accordingly, 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 first study the computations required to translate 20- to 100-line stories, expressed in simple English, into an inner story. Then, we study the computations required to use the inner story to, for example, answer questions, describe conceptual content, summarize, compare and contrast with other stories, react with cultural biases, and find useful precedents (Winston, 2011, 2012a,b).

Our 20- to 100-line stories, expressed in English, are *outer stories*. More generally, outer stories include, for example, written or spoken narratives, logs, news, reports, briefs, recipes, instructions, essays, plans, and papers. An outer story may also include or consist entirely of visual material such as found in pictures, drawings, diagrams, graphs, cartoons, and various kinds of performance. All share one property:

An outer story: Anything that produces an inner story.

In our work, however, we focus our attention on outer stories expressed in English, so when we use the word *story*, we refer to linguistically expressed outer stories. Similarly, when we refer to an inner story, we mean by default an inner story produced from linguistic input. It follows that the inner language, in which the inner story is expressed, must have certain minimal representational capabilities evident in what we say and write. We explain some of those minimal representational capabilities of any inner language when we discuss Genesis, our story understanding system, in section 2.

1.3 We formulate computational problems, posit representations, and build

Following David Marr (1982), once 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.

1.4 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 the earliest work on the Macbeth plot, is the use of *explanation rules* to account for our human tendency to look for cause: we understand that Macduff killed Macbeth because Macbeth angered Macduff, even though the causal link was 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 Eastern cultures tend to explain violence in terms of situations that lead to violent behavior, whereas people from Western cultures tend to explain violence in terms of character traits that lead to violent behavior (Morris and Peng, 1994; Awad, 2013).

1.5 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.

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 also should 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 an emerging model of basic aspects of human story understanding, not a complete model of all human story understanding. We have a decade of work mapped out just to address the obvious modeling challenges.

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 only just begun.

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? Here are some:

- 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.
- Good models support and benefit from empirical science.

We are pleased that Genesis exhibits what we think to be an impressive 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.

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

1.6 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 a 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 the Genesis Story Understanding System, 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 previous papers, Winston introduced the Genesis System, emphasized methodological steps, articulated the Strong Story Hypothesis, and saluted earlier work (2012a; 2012b), particularly the pioneering work of Schank and his colleagues and students, documented in numerous articles and books Schank (1972); Schank and Abelson (1977); Schank and Riesbeck (1981); Schank (1991).

Here, we add detail 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.

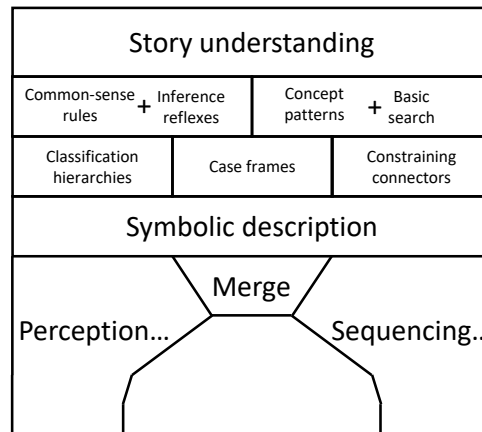


Figure 1: Genesis’s story understanding rests on a small number of surprisingly simple representations and computations. At the representation level, the Genesis system requires classification hierarchies, case frames, and constraining connectors. At the computational level, much is done with common-sense rules via inference reflexes and concept discovery via basic search. 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.

The computations embodied in Genesis, along with the enabling foundation of common-sense rules and concept patterns, constitute an evolving model of human story understanding.

2.1 Genesis reads simple, concise stories

The following short summary of *Macbeth* is the anvil on which we have hammered out many ideas:

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. Burnham Wood is a forest. Burnham Wood goes to Dunsinane. Macduff’s army attacks Macbeth’s castle. Macduff curses Macbeth. Macbeth refuses to surrender. Macduff kills Macbeth.

Shakespeare tells us a great deal about the human condition, so as we expected, the infrastructure and much of the knowledge developed to deal with *Macbeth* 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 sites showed that someone did not respect Estonia. Estonia created a center to study computer security. Estonia believed other states 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 2.8.

2.2 Genesis uses case frames extensively

Genesis uses Boris Katz’s START system to translate simple English into a collection of descriptive 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 2.6, 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 set containing just one causing element. 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 become the king, Macbeth murdered Duncan and blamed the guards*:

```
(relation means
  (sequence recipe
    (relation murder
      (entity macbeth)
      (sequence roles (function object (entity duncan))))
    (relation blamed
      (entity macbeth)
      (sequence roles (function object (entity guards)))))
  (function appear
    (relation position
      (entity macbeth)
      (sequence roles (function object (entity king)))))
```

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 by the inference reflexes 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 object whole living-thing ... bird bird-of-prey hawk
thing entity physical-entity object whole living-thing ... adult militarist hawk
```

We keep each classification sequence separate, in *classification threads*, rather than merging them into a classification tree, anticipating that we will want eventually to make use of ideas introduced by Richard Greenblatt and Lucia Vaina (1979). They note, for example, that you might want to consider a *boy* to be first a child, and then a male, or first a male, and then a child, depending on circumstances, suggesting a need for flexibility not to be found in a fixed classification tree.

2.5 Genesis uses inference reflexes to elaborate on what is written

Genesis uses actions and other story elements, together with common sense, to build an *elaboration graph*, as shown in figure 2. Elements in yellow are established by common-sense deduction rules. The story itself supplies the elements in white explicitly.

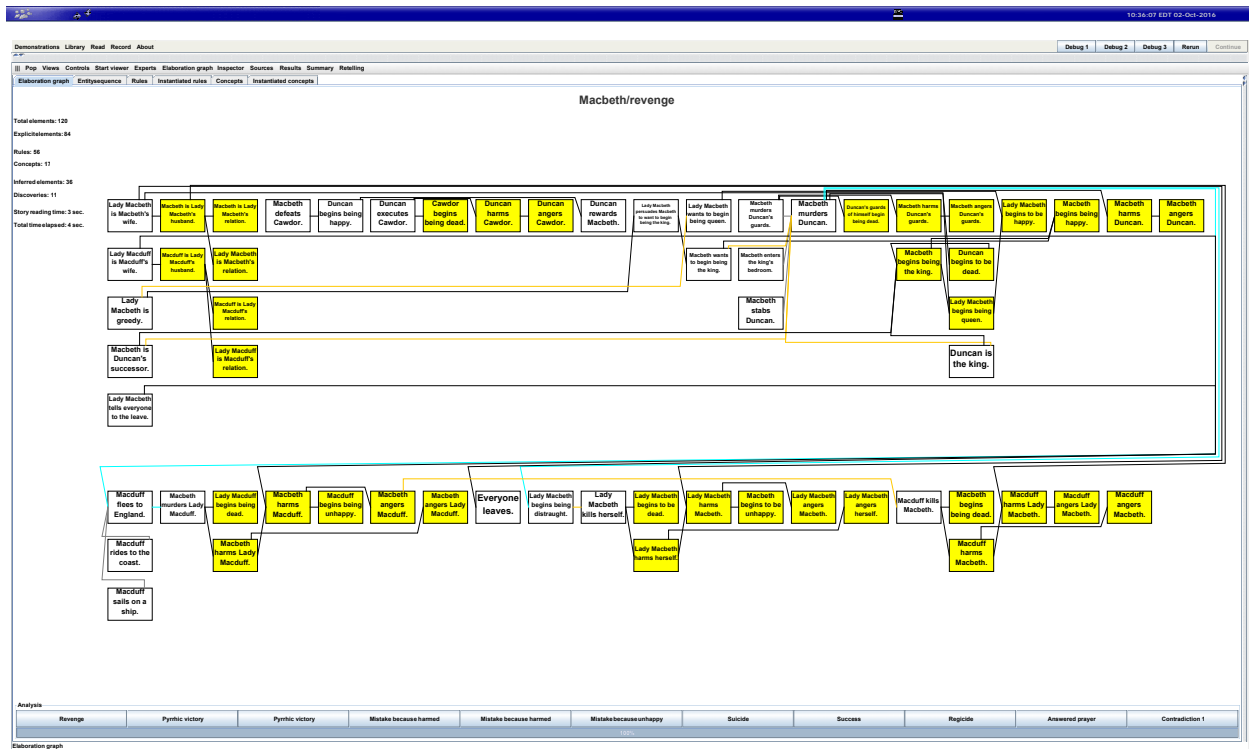


Figure 2: Genesis produces elaboration graphs as shown for a summary of Macbeth. Inference reflexes connect explicit and inferred elements of the story. Note that although the story is told as a sequence of elements, the inference reflexes form long-distance connections. (This figure is included at high resolution in the electronic version of this paper.)

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 the Genesis 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))))))
```

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 *Macbeth*, the story itself supplies no explicit reason why Macbeth murders Duncan and no deduction rule supplies a reason, but an *explanation rule* connects the murder to Macbeth's wanting to be king, Macbeth's being Duncan's successor, and Duncan's being king.

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.

We express explanation rules in English using what you can view as an idiomatic use of the word *may*, as in the follow example:

If X is king and Y wants to be king 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.

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* [Gottschall \(2012\)](#). 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.

Post hoc ergo propter hoc rules

A *post-hoc-ergo-propter-hoc rule*, also known as a *right-together rule*, is similar to an explanation rule, but post-hoc-ergo-propter-hoc rules make 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:

If X becomes Y and Z immediately becomes angry, then assume implication.

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

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): leaving out the *because* is a way of supplying no more information than is required.

Abduction rules

In reading a story, we may reach conclusions by way of cultural influence. 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.

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 deduction rule, explanation rule, or abduction rule, then the result is as if the story explicitly included *John is greedy because John is evil*.

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.

Censor rules

A seventh 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.

Inference reflexes

We have, so far, six rule types that can establish six kinds of connection, 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.

- Deduction rules, when antecedents are present, assert a conclusion and construct a deduction connection between the antecedents and the conclusion.
- Explanation rules, when both antecedents and a conclusion are present, construct a explanation connection between the antecedents and the conclusion.
- Post-hoc-ergo-propter-hoc rules, when an antecedent and a conclusion are present and next to each other, construct a proximity connection between the antecedent and the conclusion.
- Abduction rules, when a conclusion is present, assert antecedents and construct an abduction connection between the antecedents and the conclusion.
- Presumption rules assert antecedents and construct a presumption connection between the antecedents and the conclusion when a conclusion is present but no explicit cause is present, nor has any deduction, explanation, or abduction produced an explicit cause.

- Enablement rules, when a consequent is present, assert antecedents that must be true for the action to occur and construct an enablement connection between the antecedents and the consequent.
- Censor rules prevent inappropriate application of other rules.

Each such rule type was discovered 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 some story could not be properly understood. Accordingly, each rule type constitutes a *computational imperative*.

Each rule does its work the moment it can, and because each application 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.

2.6 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 read time to inference 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.

2.7 Explicit connections also contribute to basic understanding

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

Duncan became happy because Macbeth defeated Cawdor.

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

Macbeth's murdering Duncan leads to Macduff's fleeing to England.

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, Macbeth's murdering Duncan leads to Macbeth's hallucinating.

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

Macduff kills Macbeth; Macduff 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 Duncan, Macbeth stabbed Duncan.

2.8 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.

Concept discovery: 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.
```

In figure 3, Genesis notes a *Revenge* pattern because Genesis successfully searches for a sequence of causal connections between Macbeth's harming Macduff and Macduff's harming Macbeth.

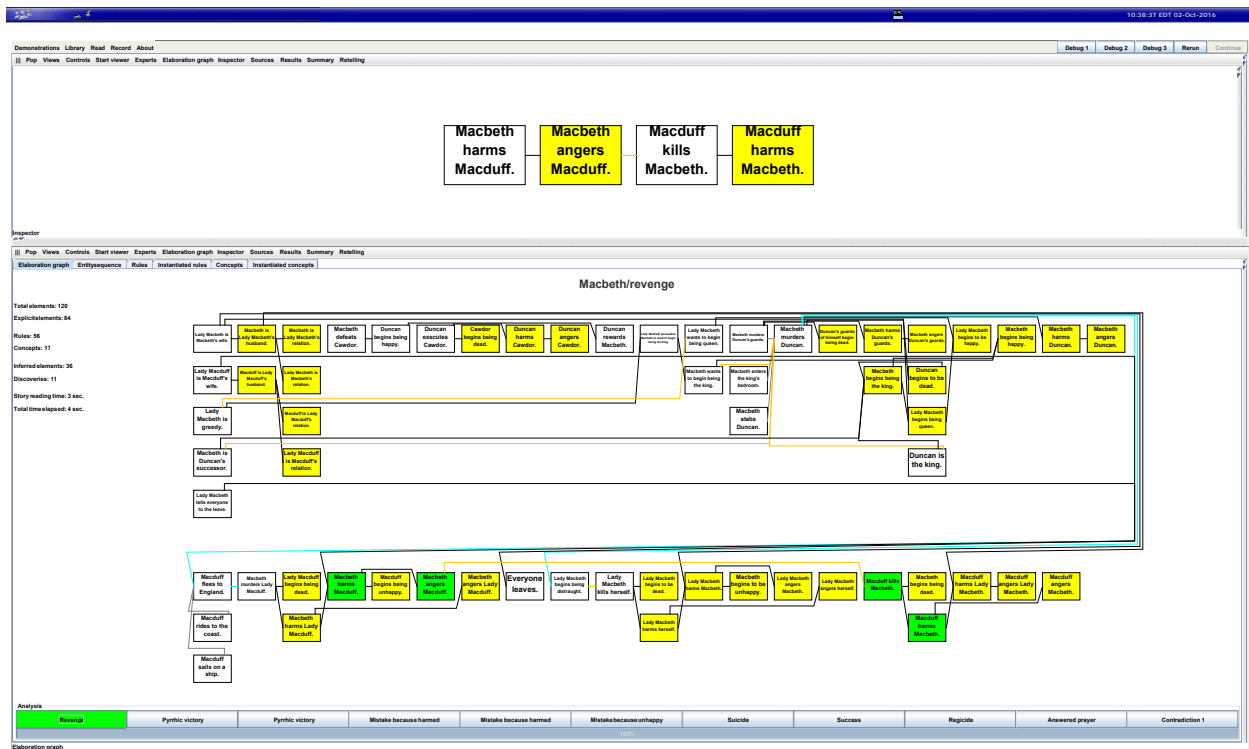


Figure 3: Genesis finds concept patterns by searching the elaboration graph. Here, Genesis highlights revenge elements in the elaboration graph. The inspector panel provides a close-up view.

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 wanting A leads to X's becoming happy.
X's wanting A leads to Y's harming X.
Y harms X after X becomes happy.
```

In figure 4, Genesis notes a *Pyrrhic victory* pattern because Macbeth's wanting to be king leads not only to becoming happy, but also leads to being harmed later.

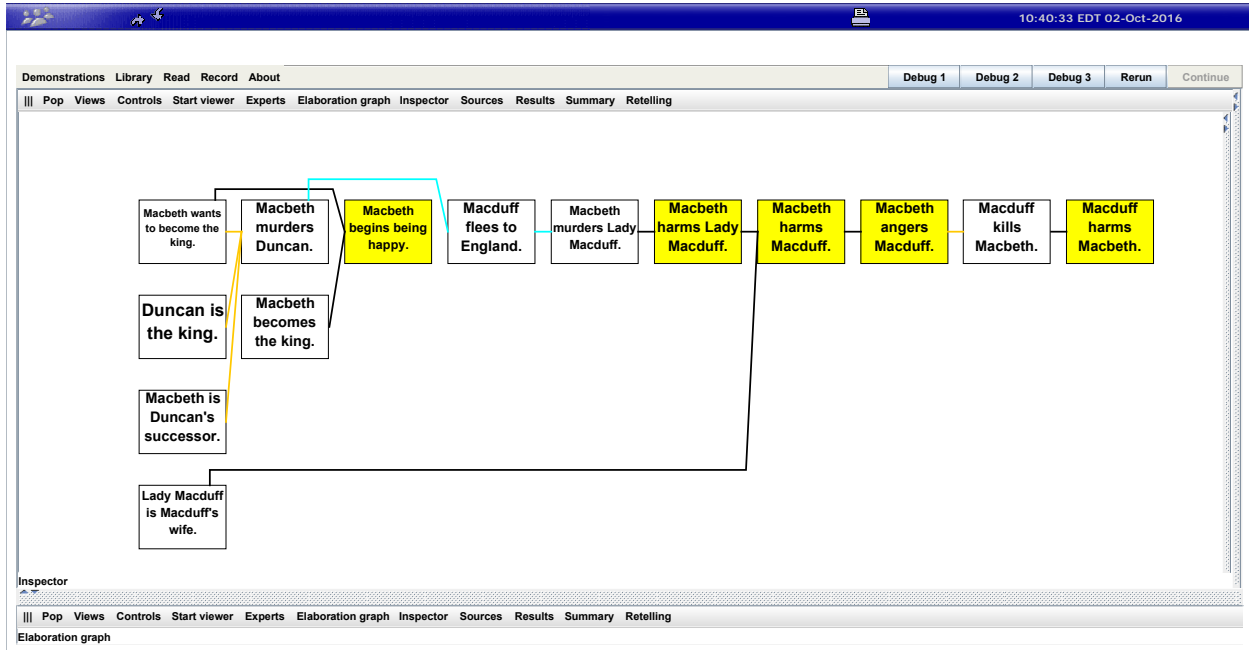


Figure 4: Genesis extracts the *Pyrrhic victory* elements from the full elaboration graph.

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 discovery 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.
```

2.9 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.

2.10 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; Genesis uses six kinds of common-sense rules to make causal connections; and Genesis uses censor rules that prevent inappropriate use of deduction, abduction, presumption, and enablement rules. Genesis uses concept patterns specifying 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 not implausibly large.

3 Genesis's simple substrate supports surprising competences

In this section, we list some of the myriad competences enabled by Genesis's small number of explicit connection types, common-sense rule types, and concept pattern types. Our purpose is to demonstrate that a simple, plausibly evolvable foundation, suffices to support myriad competences such as those shown in figure 5.




Figure 5: Basic story understanding, enabled by a small set of representational capabilities, common-sense rule types, explicit connection types, and concept pattern types, serve as a foundation for many more sophisticated story-understanding competences.

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. All the models illustrate what can be done in various dimensions; none exhaust all that can be done in any particular dimension.

Genesis answers basic questions about why and when

As shown in figure 6, Genesis answers questions on several levels, by noting the personality trait whose common-sense rule connected the target event to its antecedents, by reciting elaboration graph elements connected to the target event, and by noting how the target event is embedded in concepts.

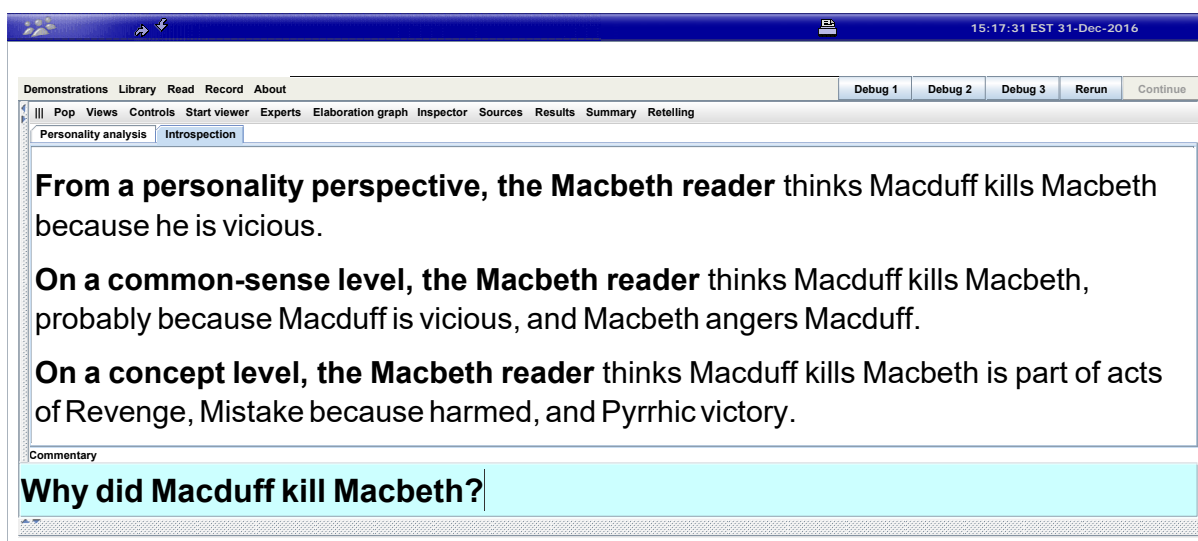


Figure 6: Genesis answers a user question, “Why did Macduff kill Macbeth,” on the personality level, the common-sense level, and the concept level.

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 an example based on the 2007 cyber war between Estonia and Russia, Genesis views the alleged actions of the Russians as misguided revenge when the story includes *Estonia is my friend*. Genesis views the same actions as teaching Estonia a lesson when the story includes *Russia is my friend* (Winston, 2012b).

Genesis’s interpretation also shifts with changes in the common-sense rules and concept patterns supplied. With one set of rules and concepts, Genesis concludes that Macbeth’s harming of Macduff causes Macduff’s killing of Macbeth; with another set, Genesis concludes that insanity causes the killing (Winston, 2012b). Thus, Genesis can exhibit either the tendency of an Eastern reader to reason situationally, or the tendency of a Western reader to reason dispositionally, empirically established by Michael W. Morris and Kaiping Peng (1994).

Genesis models personality traits

Genesis notes what various sorts of people do, which 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 Macbeth story that Macduff assaults someone, an act Genesis has recorded as indicative of vicious people, leading Genesis to consider Macduff to be vicious. Then, whenever Macduff is involved in an action, common-sense rules associated with viciousness are added to those generally used.

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 7, the potential for revenge, misguided retaliation, and mistake are noted early in one version of the Estonia–Russia cyber war. All three eventually ensue, as shown in figure 8.

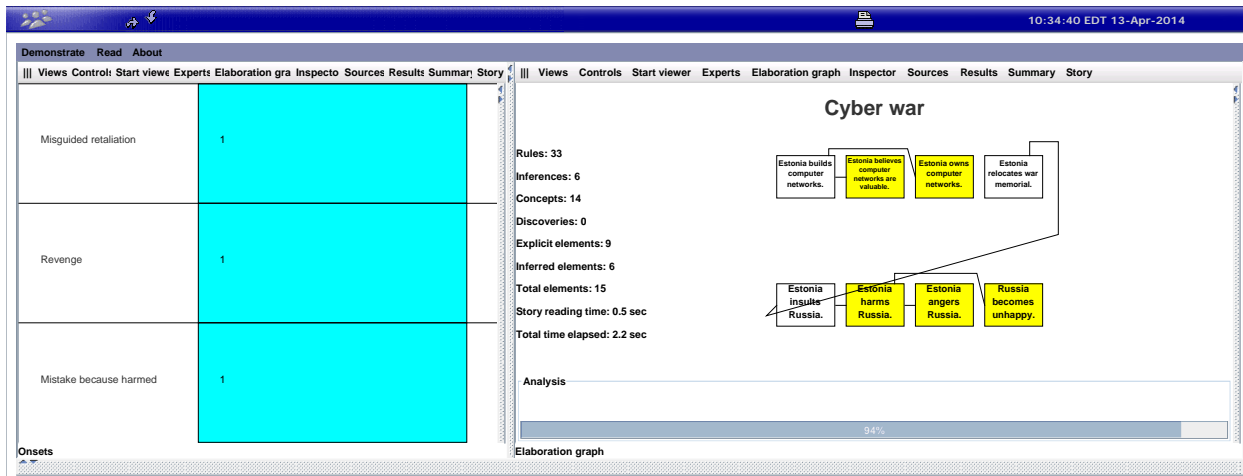


Figure 7: Genesis notes the onset of three possible concepts midstream in the Estonia–Russia cyber-war story.

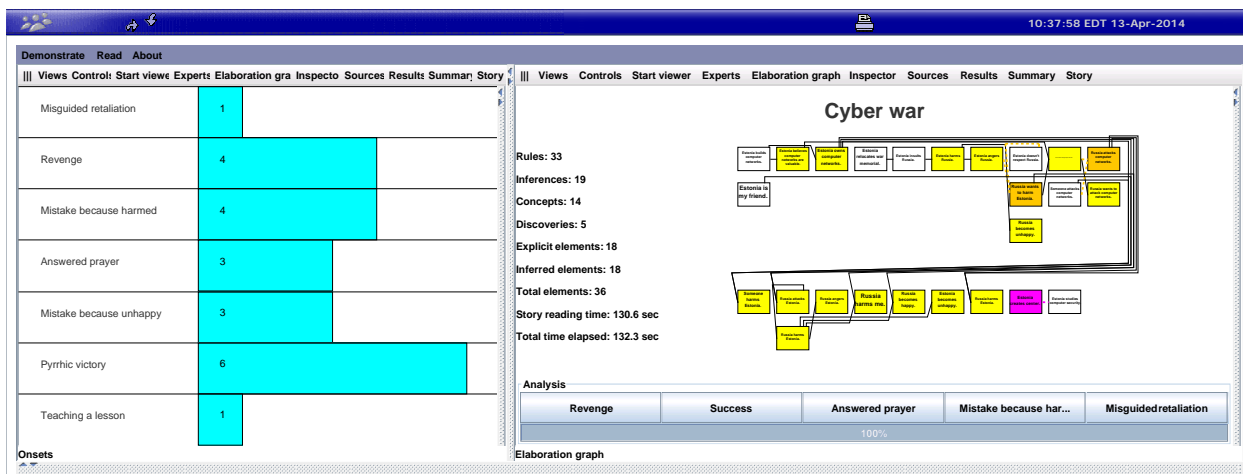


Figure 8: Genesis concludes that all three of the anticipated concepts eventually become realized.

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 9, 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.

Story	Element 0	Element 1	Element 2	Element 3	Element 4	Element 5	Element 6	Element 7	Element 8
Story B Score: 6.0	Israelis know that Egyptians are preparing to attack them.	Egyptians are preparing to attack Israel.	Israelis defeat Egyptians.	Israelis know to defeat Egyptians.	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.
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.	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 9: Genesis aligns elements in two wars and fills gaps in each using the other.

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).

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 10 are on pairs of short descriptions of conflicts (Krakauer, 2012).

Genesis models question-driven interpretation

After reading a story, a question may stimulate further analysis and expose new conclusions. The example here is from an Eastern-Western story understanding demonstration (Morris and Peng, 1994; Awad, 2013).

In the story, a student murders a professor and another student. Genesis, modeling an Eastern reader, has no opinion on why Lu killed Shan until asked if it was because America is individualistic. Then, as shown in figure 11, having been asked a question, Genesis recalls that the question's antecedent is something that the reader believes, which leads to adding that recalled belief to the story, with consequences that connect the inserted belief to the murder. Genesis affirms that Lu killed Shan because America is individualistic.

Another version of Genesis, modeling a Western reader, recalls no such belief, so fails to insert *America is individualistic* 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.




Figure 10: Genesis performs concept-based similarity measurements. Concept-based measurements are shown above and word-based similarity measurements below. White means most similar. Note that the similarity conclusions reached using concepts are different from those reached using only words.

Genesis develops summaries around conceptual content

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

In figure 12, for example, limiting the telling of *Macbeth* to those elements connected to the Pyrrhic victory concept compresses the summary provided into a shorter summary by about 7:1 (Winston, 2015).

Genesis tells and persuades 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 13, Genesis supplies principles to a reader that knows very little in the beginning, but is taught that, for example, you become king if the present king dies and you are his successor.

Similarly, Genesis can tailor what is said to shape reader opinion. In figure 14, 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).

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). In the following, Fay’s system has reused elements from Shakespearean tragedies and war stories:

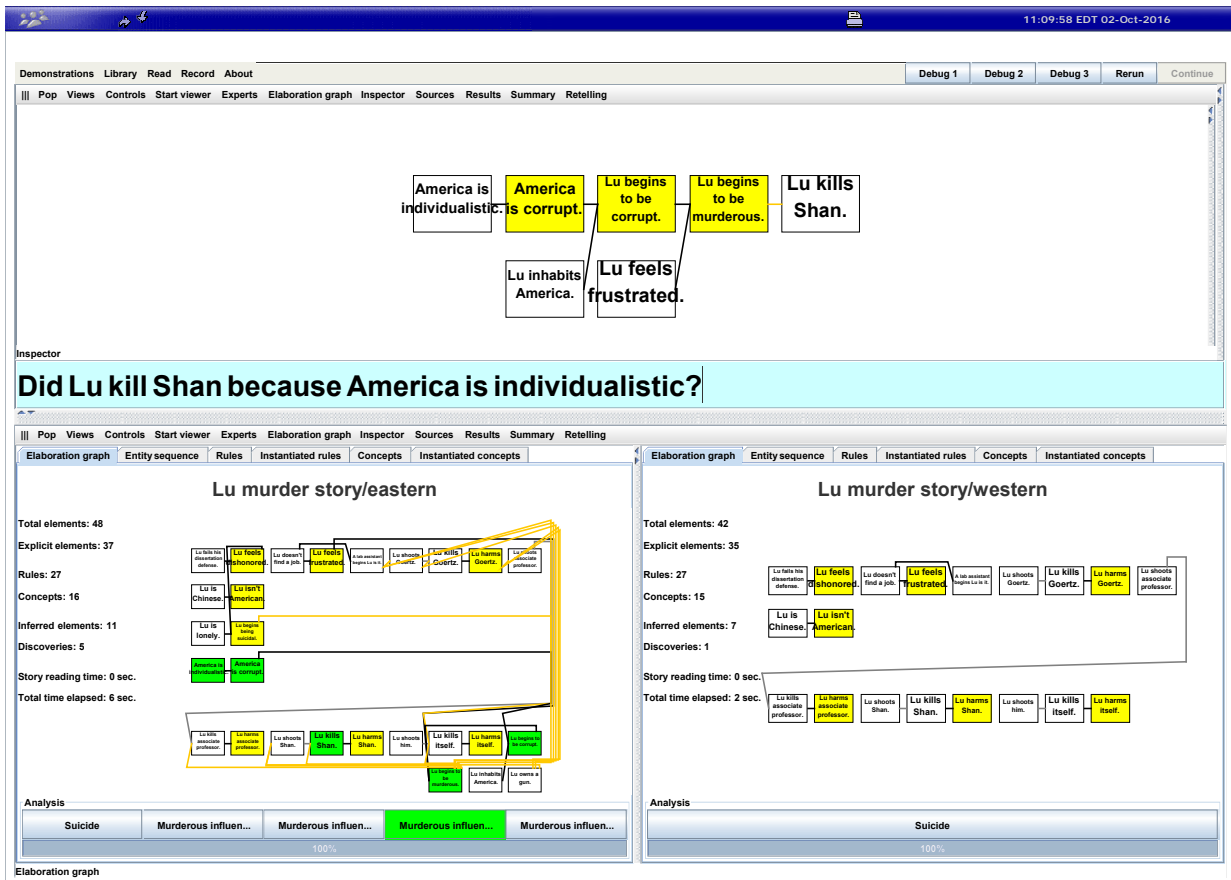


Figure 11: Genesis interprets the Shan murder. The basic interpretation does not connect the murder with America until after a question is asked and a belief is inserted by the Genesis version that models an Eastern reader. Then, a *Murderous influence* connection is noted by the Genesis model of an Eastern reader, but not for the Genesis model of a Western reader.

Greinia and Astilir

Greinia is a country. Astilir is a city in Greinia. Angelina is from Astilir. 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. Astilir rebels. Astilir attacks Greinia. Malcolm scares away Astilir's forces. Astilir attacks Malcolm. Astilir kills Malcolm. Astilir defeats Greinia. Astilir becomes independent.

In the following, 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.

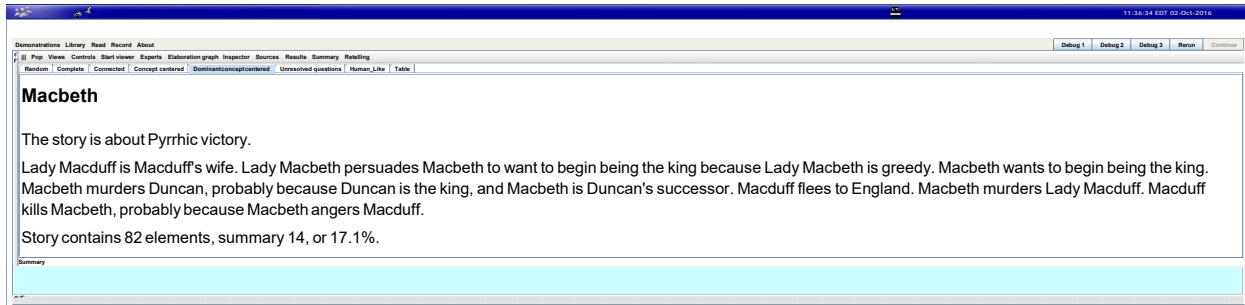


Figure 12: Genesis summaries *Macbeth* by keeping only the explicit story elements that are connected to the Pyrrhic victory concept.

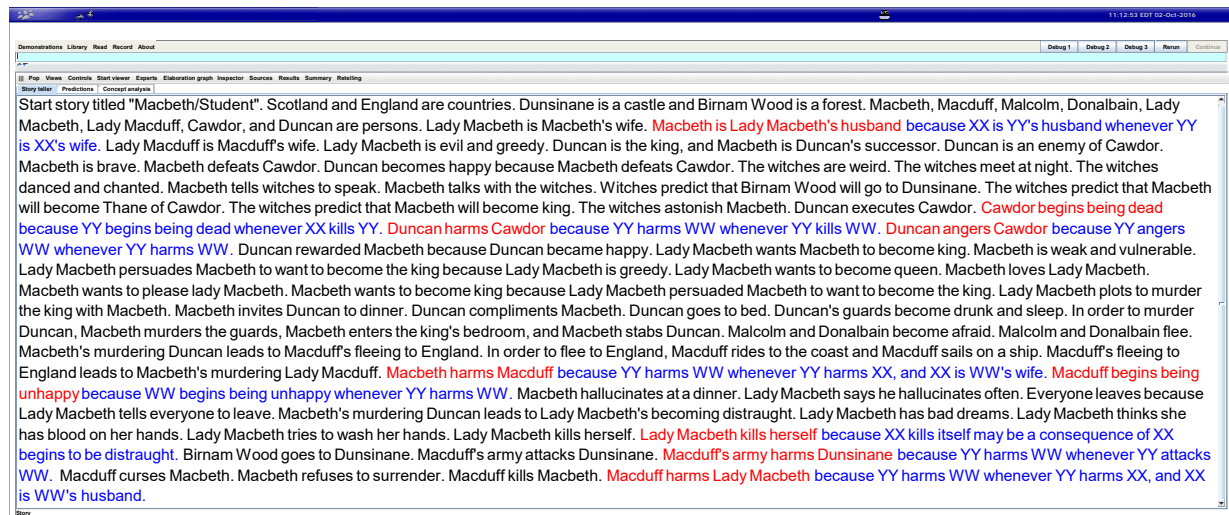


Figure 13: Genesis uses a reader model to determine what and how much to say in retelling *Macbeth*. Here, Genesis says a lot, because Genesis's model of the reader suggests that the reader does not know much.

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 Misérables*, 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 15, for example, Genesis answers a question about what Javert believes.

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 system:

Why does Jean Valjean disagree with Inspector Javert?
Inspector Javert and Jean Valjean disagree about "Jean Valjean is criminal".

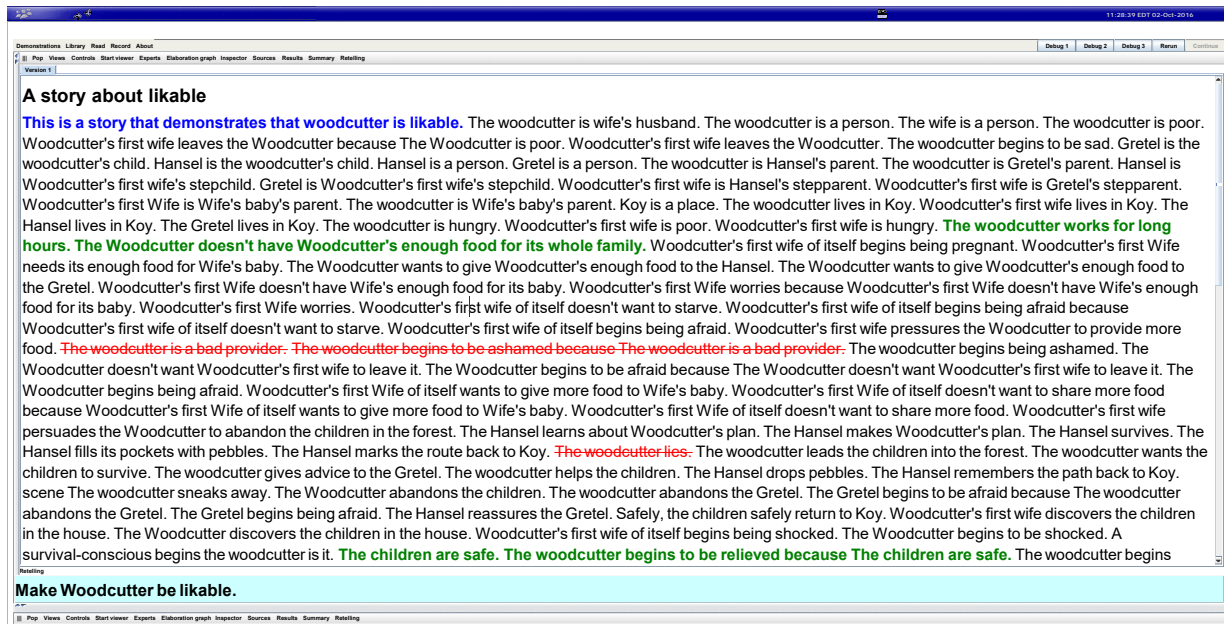


Figure 14: Genesis uses a reader model to determine what and how much to say so as to shape the reader's opinion in this first part of a retelling of Hansel and Gretel. The good is emphasized; the bad is struck out.

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 commonsense (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).

Genesis reasons about hypothetical possibility

When asked a hypothetical question about a story, Genesis need only reinterpret the story with indicated additions or deletions.

Suppose, for example, that you ask Genesis how its analysis of a story of a break-in would change if the intruder had no weapon. To answer, Genesis removes the indicated element from the story, re-analyzes the story, and produces a summary describing the differences at both a fine-grained story-element level and at an abstract, conceptual level as shown in figure 16.

Working with the Estonia–Russia story discussed in section 2.1, Genesis first finds the concept *Aggression of a bully* because of an added statement, *I am a friend of Estonia*. Then asked, *What would happen if I am not a friend of Estonia?* Genesis concludes that the interpretation should be *Teaching a lesson*.

Working with a version of Macbeth story shown in section 2.1, several hypothetical questions all lead to the conclusion that Macduff no longer kills Macbeth:

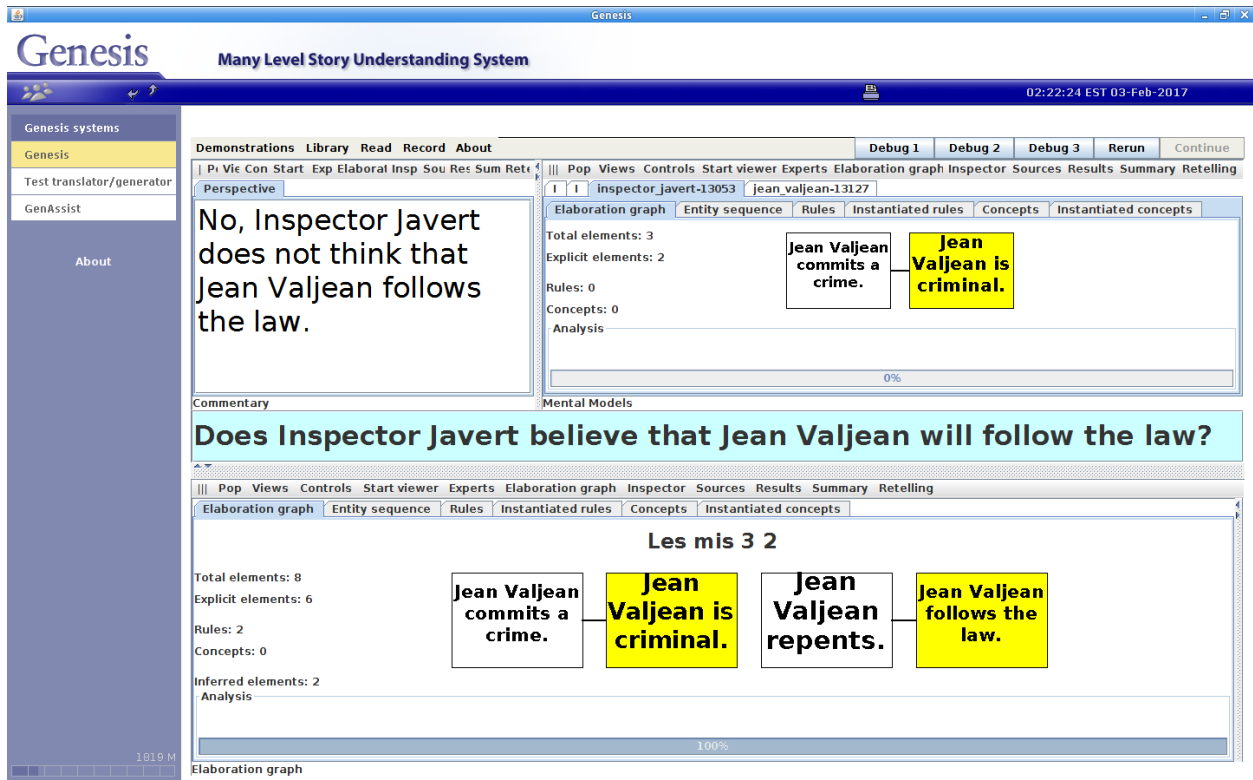


Figure 15: Genesis answers comprehension questions about characters' perspectives. Here, Genesis observes that the policeman will not believe that this criminal follows the law because the policeman was not present when the criminal repented.

- What would happen if Macduff were not an enemy of Cawdor?
- What would happen if Lady Macbeth were not greedy?
- What would happen if Macbeth did not want to be king?
- What would happen if Macbeth did not murder Lady Macduff?

These examples, from legal-reasoning, cyberwar, and literature, illustrate the power and wide applicability of hypothetical reasoning based on story understanding. 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.

3.2 Aspects of many other competences will be demonstrated

Various forms of learning populate our list of scaling-up challenges. Right now, we supply all the common-sense rules and concept patterns in English, as a parent or teacher would when explaining a story to a child or student. Of course, we humans formulate some rules and concepts independently, not only by asking questions but also by self discovery from experience and the surrogate experience provided by reading.

How much do we learn by self discovery? How can we model the self-discovery process? So far, we have ideas, but no complete results, except for Mark Finlayson's (2012) groundbreaking work on modeling Vladimir Propp's extraction, from Russian folk tales, of what we would call concept patterns.

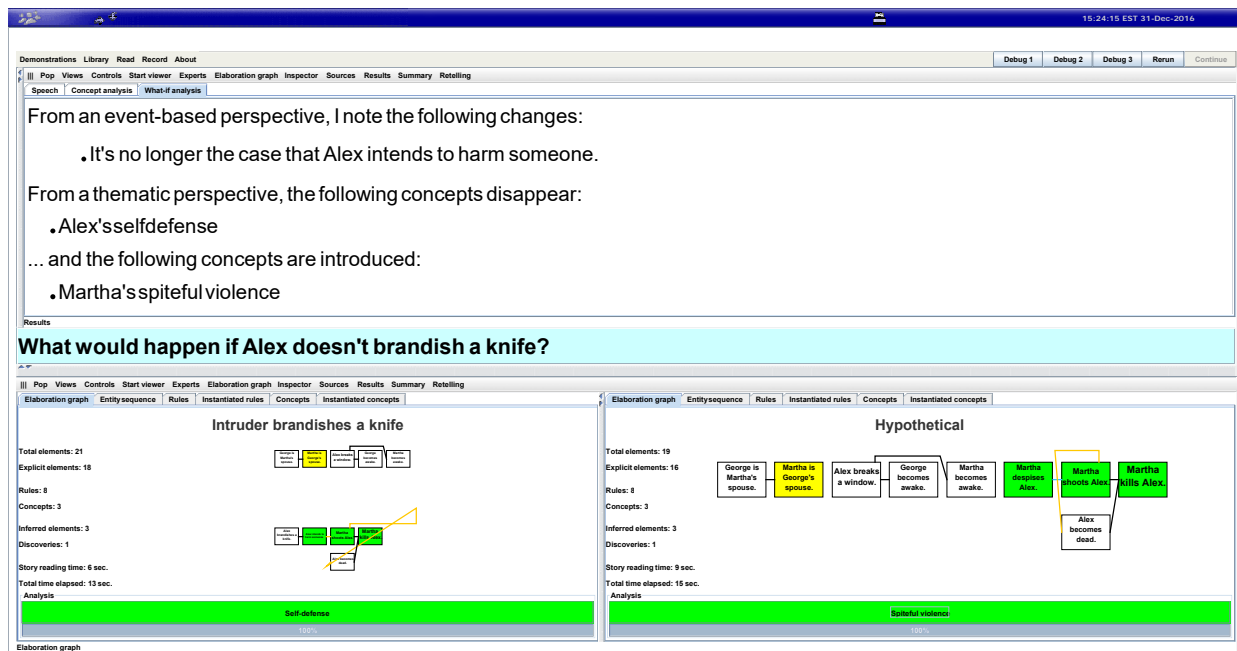


Figure 16: In the knife-brandishing version of a story raising a legal question, the elaboration graph indicates the brandishing is connected to the killing, suggesting self defense. In the hypothetical version, with the brandishing removed, Genesis presumes the explanation for killing has to do with despising, suggesting guilt. At the concept level, *Self defense* becomes *Spiteful vengeance*.

Another scaling-up challenge is to find ways to 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 our friends focused on internalization/externalization will take on those challenges for us.

3.3 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. Many more such steps are underway; many others are obvious and ready for investigation.

4 Work on story understanding is fundamentally multidisciplinary

Our purpose in this section is to suggest that story understanding is a multidisciplinary enterprise. First, we identify representative work that we have found to be especially influential in Artificial Intelligence and in empirical science; next, we explain how our work suggests challenges for those who work on neural nets; and finally, we describe what our work has to offer other fields.

4.1 We are inspired and encouraged by Minsky’s work in Artificial Intelligence

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).

4.2 We are inspired and encouraged by results from empirical science

We think our views are consistent with results from empirical science that we consider to be particularly suggestive. Elizabeth Spelke and Linda Hermer-Vazquez note that verbal shadowing—subjects repeat what they hear as they hear it—interferes with combining geometric and nongeometric features (1999): we conjecture that verbal shadowing jams the inner-language story-processing competence such that adult-human subjects cannot even tell themselves a two-step story about how to find which of four corner baskets holds an item of interest (go to blue wall, turn left).

Likewise, we are impressed by the work of Evelina Fedorenko and Rosemary Varley, who note that subjects with totally nonfunctional linguistic cortex still play chess and show other signs of thinking (2016): we conjecture that a relatively intact inner-language story-processing competence still functions even though the externalization/internalization cortex is gone.

4.3 Our work suggests problems for empirical 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:

- 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.

4.4 Our work 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.

4.5 Our work has much to offer to many fields

- For *computer science and artificial intelligence*, we offer a prototype for systems with humanlike intelligence that communicate with humans in human terms.
- For *psychology*, we offer theories described in computational language—theories that are concretely implemented, tweakable, and testable by virtue of being computational.
- For *neuroscience*, we offer suggestions about what sorts of computations to look for in neural hardware.
- For *design*, we offer a starting point for thinking about how symbolic and visual thinking stimulate each other in creative action and retrospective critique.
- For *philosophy*, we offer a rigorous methodology for exploring and expressing questions about the mind. Concepts like *virtual machines*, for example, can advance long-standing philosophical questions about functionalism and supervenience
- For *linguistics*, we offer 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*, we offer a characterization of what makes us quintessentially human.
- For *literary studies*, we offer a way of thinking about how stories become structured, coherent, surprising, and memorable.
- For *education*, we offer 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*, we offer 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*, we offer ways to mitigate some of the dangers that concern anxious futurists. By developing self-reflective architectures, we offer machines that can explain themselves—our only hope for safety as we come to depend more and more on our intelligent artifacts. And we offer the possibility of unprecedented prosperity, through machines that can think and plan and aid human users.

4.6 Summary

We have been much inspired and encouraged by work in multiple fields. We are inspired and encouraged, for example, by Marvin Minsky’s work on levels of reasoning and consciousness and by Elisabeth Spelke’s work on the cognitive competences that separate us humans from other animals.

We believe our work has much to offer empirical 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 have suggested 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 offer a way of thinking and the potential for new and exciting tools.

5 Contributions

In this manifesto, we have explained why we are passionate about developing computational models of our human ability to construct, tell, and understand stories. In particular, we have:

- Argued that the ability to assemble complex, highly nested symbolic descriptions enabled the story competences that make us human. 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 the *computational imperative principle* and the role of computational imperatives in developing scientific, computational accounts of intelligence.
- Exhibited many kinds of common-sense rules and concept patterns, showing how they are put to work in *inference reflexes* and *concept discovery* so as to support models of many story-understanding competences.
- Enumerated many story-understanding achievements, from models of question answering to story-grounded hypothetical reasoning, and identified current challenges.
- Discussed how our work on the Genesis model is aligned with empirical science, informed by empirical science, and can contribute to empirical science.
- Suggested how success in story understanding benefits 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.

6 Acknowledgements

The work described in the paper was supported, in part, by the Air Force Office of Scientific Research. Other sponsors of work on aspects of the Genesis Group's work include the National Science Foundation and the Defense Advanced Research Projects Agency.

References

- Hiba Awad. Culturally based story understanding. Master's thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 2013.
- Robert C. Berwick and Noam Chomsky. *Why Only Us*. MIT Press, Cambridge, MA, 2016.
- 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.
- Charles J. Fillmore. *The Case for Case*, pages 1–88. Holt, Rinehart, and Winston, 1968.
- Mark A. Finlayson. *Learning Narrative Structure from Annotated Folktales*. PhD thesis, Electrical Engineering and Computer Science Department, MIT, Cambridge, MA, 2012.
- 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 Publications, London, 2007.
- 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.
- Linda Hermer-Vazquez, Elizabeth Spelke, and Alla S. Katsnelson. Sources of flexibility in human cognition: Dual-task studies of space and language. *Cognitive Psychology*, 39(1):3–36, 1999.
- 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.
- David Marr. Artificial intelligence—a personal view. *Artificial Intelligence*, 9:37–48, 1977.
- David Marr. *Vision*. W. H. Freeman, San Francisco, CA, 1982.
- Marvin Minsky. *The Society of Mind*. Simon and Schuster, New York, NY, 1988.
- Marvin Minsky. *The Emotion Machine*. Simon and Schuster, New York, NY, 2006.
- 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.
- 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.
- Livia Polanyi. *Telling the American Story: A structural and Cultural Analysis of Conversational Storytelling*. MIT Press, Cambridge, MA, 1989.
- 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, 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. Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology*, 3(4):552–631, 1972.
- Roger C. Schank. *Tell Me a Story: A New Look at Real and Artificial Intelligence*. Northwestern University Press, Evanston, IL, 1991.
- 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.
- 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*. 2012.
- Shimon Ullman. *High-Level Vision*. MIT Press, Cambridge, MA, 1996.
- Lucia M. Vaina and Richard Greenblatt. The use of thread memory in amnesic aphasia and concept learning. Working Paper 195, Artificial Intelligence Laboratory, MIT, Cambridge, MA, 1979.
- Patrick Henry Winston. The strong story hypothesis and the directed perception hypothesis. In Pat Langley, editor, *Technical Report FS-11-01, Papers from the AAAI Fall Symposium*, pages 345–352, Menlo Park, CA, 2011. AAAI Press.
- Patrick Henry Winston. The next 50 years: a personal view. *Biologically Inspired Cognitive Architectures*, 1, 2012.
- Patrick Henry Winston. The right way. *Advances in Cognitive Systems*, 1:23–36, 2012.
- Patrick Henry Winston. Model-based story summary. In Mark A. Finlayson, Ben Miller, and Remi Ronfard, editors, *Proceedings of the 6th Workshop on Computational Models of Narrative(CMN 2015)*, volume 45. Oasics, 2015.
- 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.