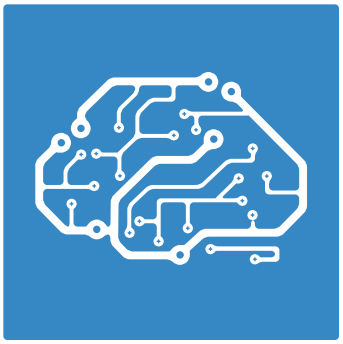

Center for Brains, Minds & Machines

CBMM Memo No. 019 **June 10, 2014**

The Genesis Story Understanding and Story Telling System
A 21st Century Step toward Artificial Intelligence

by

Patrick Henry Winston

Abstract: Story understanding is an important differentiator of human intelligence, perhaps the most important differentiator. The Genesis system was built to model and explore aspects of story understanding using simply expressed, 20-100 sentence stories drawn from sources ranging from fairy tales to Shakespeare's plays. I describe Genesis at work as it reflects on its reading, searching for concepts, reads stories with controllable allegiances and cultural biases, models personality traits, answers basic questions about why and when, notes concept onsets, anticipating trouble, calculates similarity using concepts, models question-driven interpretation, aligns similar stories for analogical reasoning, develops summaries, and tells and persuades using a reader model. I conclude with thoughts on how Genesis would describe people in pictures and video, thus engaging with the CBMM challenge problem.

This work was supported, in part, by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF - 1231216.

Vision: a distinguishing difference and a shared substrate

Marvin Minsky published *Steps toward Artificial Intelligence* a little more than 50 years ago (1961). Now, half a century later, Deep Blue, Watson, and Siri amaze everyone and stand as testaments to engineering skill, but the applications dreamed of in 1961 remain far from realized:

- No language system has the common sense required to understand what it is reading and the implications that flow from common sense understanding.
- No vision system understands the visual world well enough to report instances of more than a handful of the 48 actions cited in DARPA's 2010 Mind's Eye BAA.

Why are there no such systems? I think it is because we remain ignorant of how our species is different from others, and we remain ignorant of how the differences are enabled by a shared substrate.

In this white paper, I focus on story understanding, which I believe is the key difference between us and all other primates, living and extinct. I catalog our work on story-understanding tools, describing some of the capabilities of our Genesis story-understanding system.

Steps: Specify behavior, develop representations, build Genesis

It is evident that we learn a great deal about life from stories, which include fairy and folk tales, religious parables, ethnic narratives, entertainment, news, history, literature, and experience. It is also evident that we learn a great deal professionally from case studies in law, business, medicine, defense, diplomacy, science, and engineering.

To understand how we understand and learn from stories, I believe we must first specify the behavior we want to model, then develop a suite of constraint exposing representations, then build our Genesis story understanding system so as to test our ideas. I believe success will lead to understanding human intelligence and to paradigm-shifting engineering practice built on that understanding.

What Genesis models

I have elaborated on our perspective in a trilogy: The Strong Story Hypothesis and the Directed Perception Hypothesis (2011) introduces two key hypotheses about human intelligence and ties story telling to directed perception; The Right Way (2012b) emphasizes mythological steps; and The next 50 years: A personal view (2012a) focuses on asking the right questions about how we are different from other primates and how we are the same.

My purpose here is to add to those previous papers an account of what the Genesis system now models by way of reading, deliberating, reflecting, cultural bias, personality understanding, question answering, onset detection, similarity measurement, similarity based retrieval, question driven interpretation, analogical interpretation, reader aware story telling, persuasion, and summary.

Genesis reads stories, common sense rules, and concept descriptions

We of the Genesis Group spend a lot of time with a short summary of *Macbeth*. Shakespeare, in general, tells us about the human condition, his plays constitute a good anvil for hammering out story understanding 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. Birham 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. The end.

I noted early on that what works for *Macbeth* works also for other kinds of conflict. All of the infrastructure and much of the knowledge developed to deal with *Macbeth* transferred over to our work on the Estonia-Russia cyber war of 2007.

Estonia built Estonia's computer networks. Estonia insulted Russia because Estonia relocated a war memorial. Russia wanted to harm Estonia. Someone attacked Estonia's computer networks after Estonia harmed Russia. Russia 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. The end.

The Estonia-Russia cyber war is much like other conflicts. Dealing, for example, with the Congo civil war in a retrieval experiment introduced no challenges.

Genesis deploys common sense rules to develop basic understanding

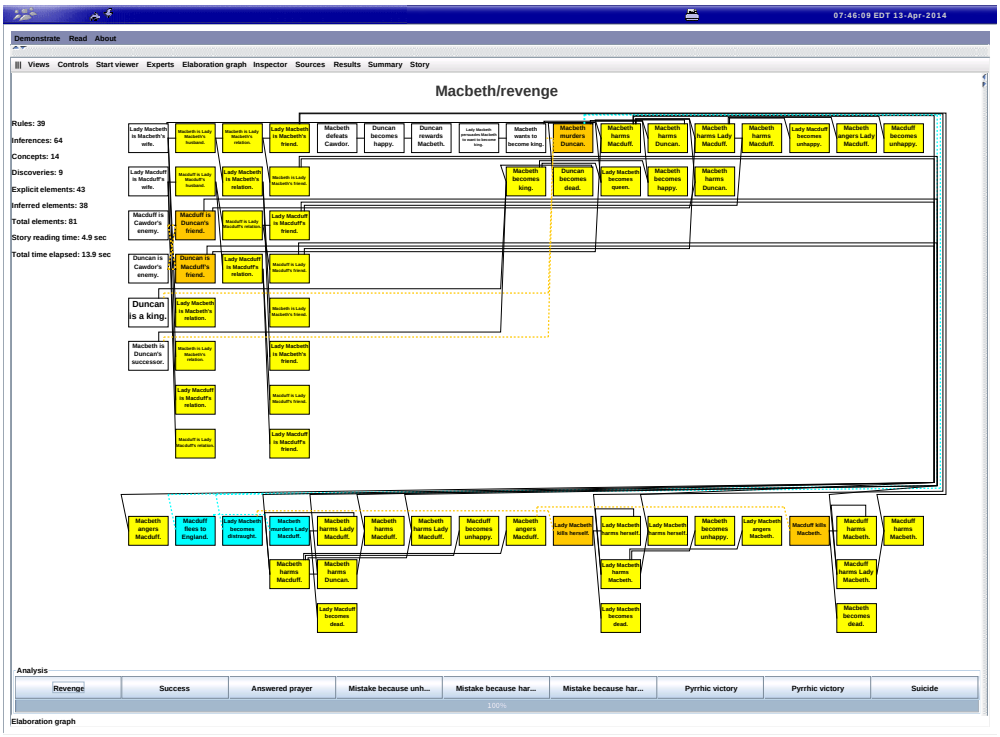

Figure 1: Genesis produces elaboration graphs, as shown for a summary of Macbeth. Common sense rules connect explicit and inferred elements of the story. (This figure is included at high resolution in the electronic version of this document.)

Genesis uses Boris Katz's START (1997) system to translate Genesis English into an inner language of relations and events. Genesis then uses common sense to build an *elaboration graph* as shown in figure 1. Elements in yellow are established by *inference rules* as indicated by black connecting lines. The story itself supplies the elements in white, orange, and blue.

Explanation rules tie elements in orange to other elements. 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

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

Leads-to expressions connect elements in blue to other elements. These are supplied by expressions in the story, such as *Macbeth's murdering Duncan leads to Macduff's fleeing to England*. Such expressions indicate when two elements are known to be connected but the exact causal path is not known, or at least not supplied.

Genesis reflects on its reading, searching for concepts

Once Genesis builds the elaboration graph, Genesis uses ordinary search to find instances of concept patterns. In figure 2, Genesis notes a *revenge* pattern because Macbeth's harming Macduff leads to Macduff's harming Macbeth. In figure 3, Genesis notes a *Pyrrhic victory* pattern because Macbeth's actions at first make him happy, but then lead to his own harm.

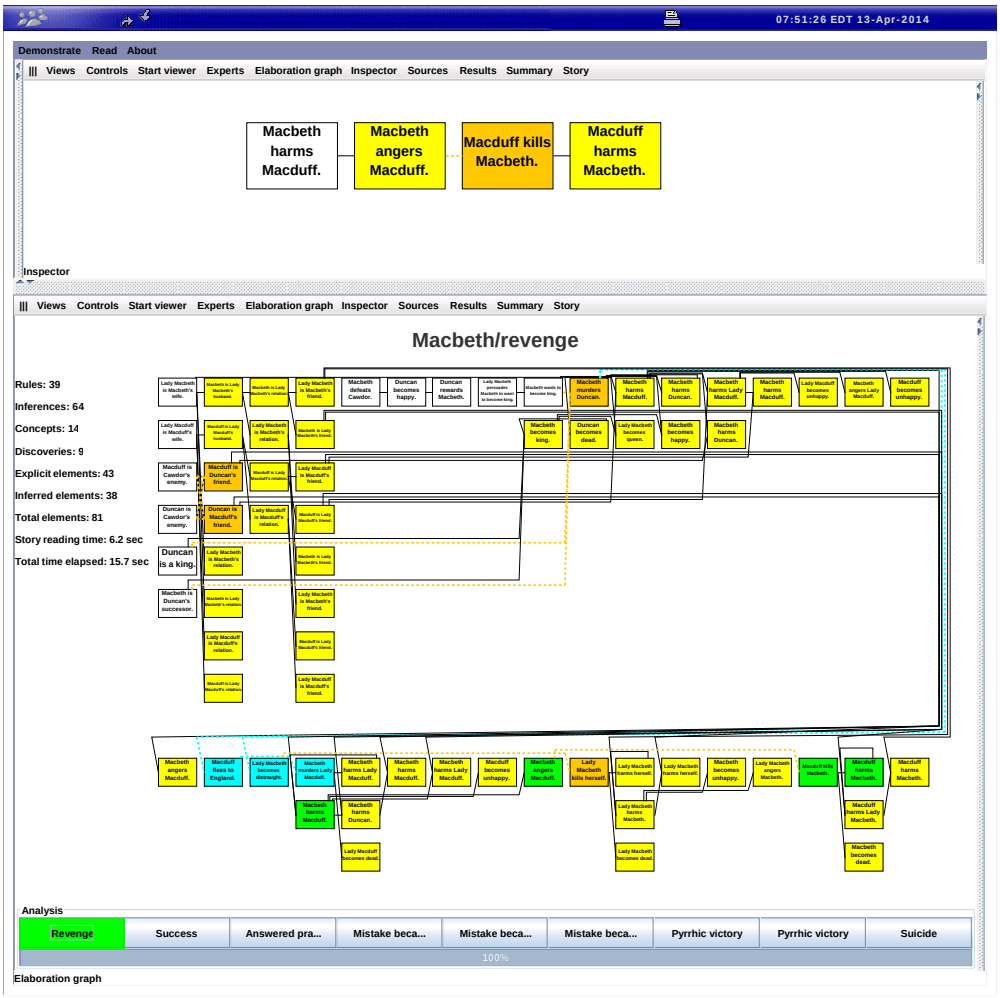

Figure 2: Genesis finds concept patterns by searching the elaboration graph. Here, Genesis shows revenge elements in green in the Elaboration graph panel provides a close-up view in the Inspector panel.

Once Genesis builds the elaboration graph, Genesis uses ordinary search to find instances of concept patterns. In figure 2, Genesis notes a *revenge* pattern because Macbeth's harming Macduff leads to Macduff's harming Macbeth. In figure 3, Genesis notes a *Pyrrhic victory* pattern because Macbeth's actions at first make him happy, but then lead to his own harm.

Our first work on concept patterns appeared in the MEng thesis of David Nackoul (2010).

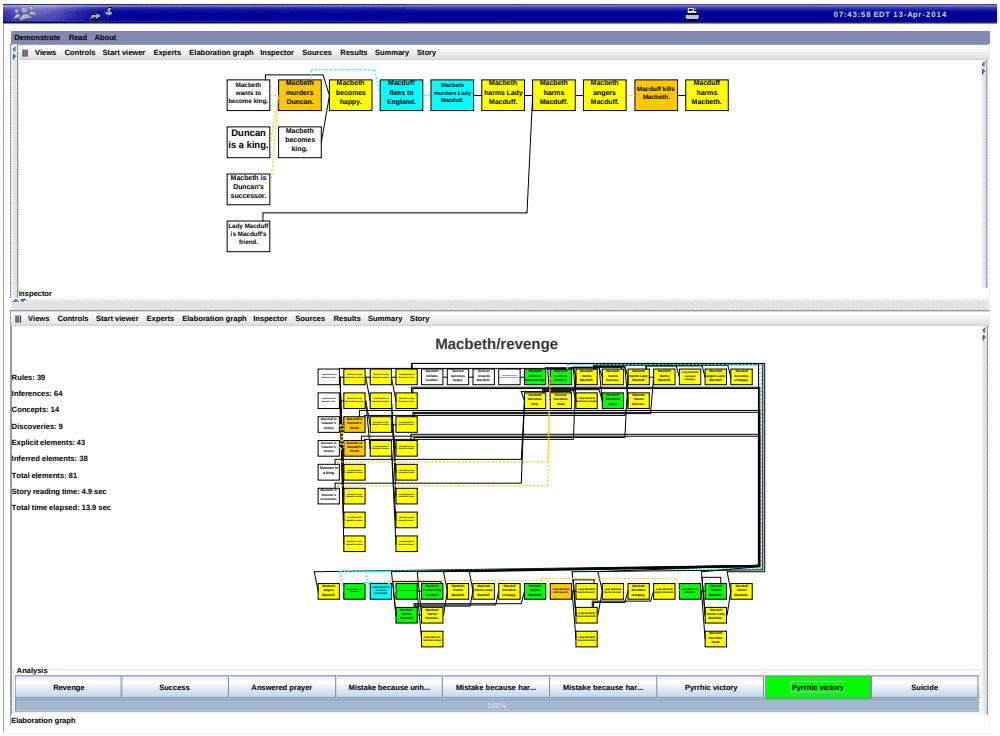

Figure 3: Genesis extracts the pyrrhic victory elements from the full elaboration graph.

Genesis reads stories with controllable allegiances and cultural biases

Genesis's interpretation depends on the common sense rules, concept patterns supplied, and biases of the reader. In figure 4, the 2007 cyber war between Estonia and Russia is viewed as misguided revenge buy a reader friendly to Estonia; the same cyber war is viewed as teaching Estonia a lesson by a reader friendly to Russia.

Similarly, in figure 5, an Eastern reader of Macbeth is more likely than a Western reader to see Macduff's killing of Macbeth situationally, as a consequence of the situation Macduff is in; a Western reader is more like to see the same killing dispositionally, as a consequence of insanity.

The most ambitious examples of cultural influence appear in the MEng thesis of Wolfgang Victor Yarlott (2014), who demonstrated Genesis at work on selected folktales from Native American Crow culture.

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.

In figure 6, Genesis notes early in the Macbeth story that Macduff assaults someone, a characteristic of vicious people, leading Genesis to consider Macduff to be vicious. Then, thinking that Macduff is vicious, Genesis explains Macduff's killing of Macbeth as a consequence of anger and Macduff's vicious nature.

Our first work on personality traits appeared in the MEng thesis of Susan Song (2012).

Genesis answers basic questions about why and when

Genesis answers questions on various levels. As shown in figure 7, Genesis answers questions about cyber war by reciting elaboration graph elements connected to the target event and by noting how target events are embedded in concepts. Additionally, as shown in figure 8, Genesis answers questions using personality traits associated with the target event.

Genesis notes concept onsets, anticipating 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 9, the potential for revenge, misguided retaliation, and mistake are noted

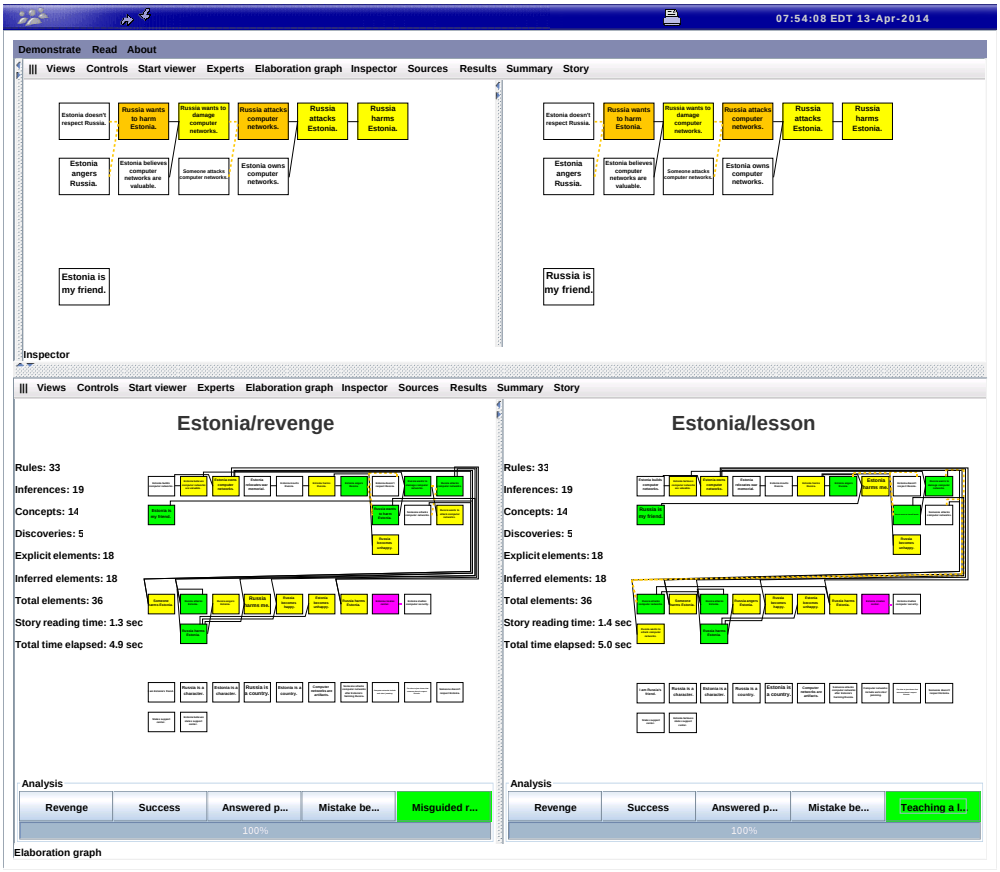

Figure 4: Genesis views the 2007 cyber war between Estonia and Russia from the perspective of a friend of Estonia on the left and from the perspective of a friend of Russia on the right. One sees misguided revenge; the other, teaching a lesson.

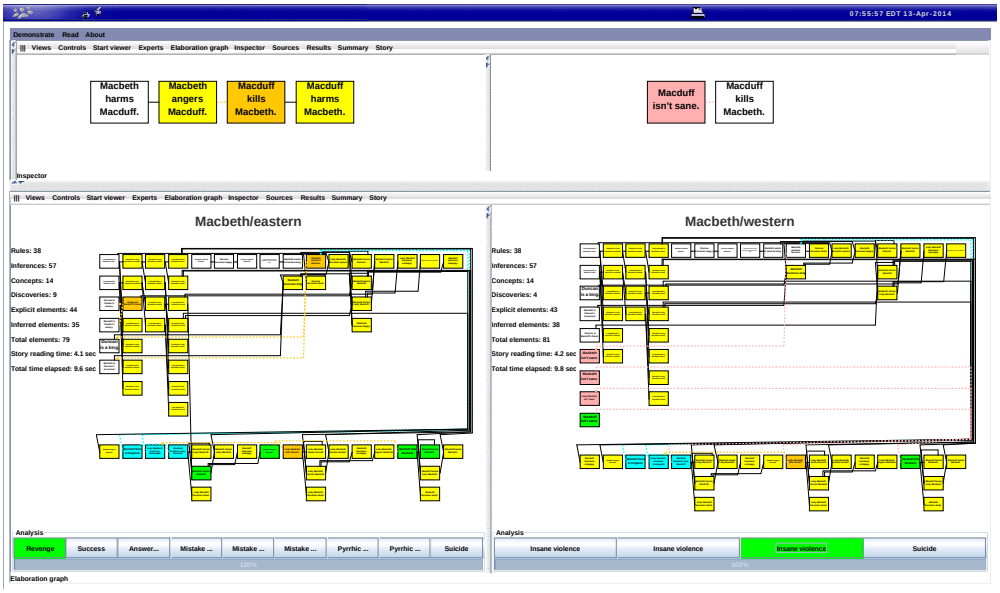

Figure 5: Genesis views the killing of Macbeth from Eastern and Western points of view. One the left, the Eastern view, the killing is situational; on the right, the Western view, dispositional.

early. All three eventually ensue, as shown in figure 10.

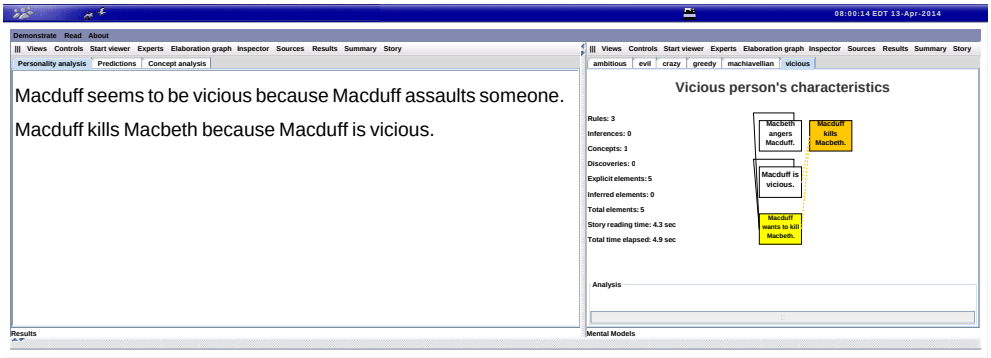

Figure 6: Genesis notes early in the Macbeth story that Macduff has done a vicious act, which leads Genesis to use rules associated with visciousness further along in understanding the story.

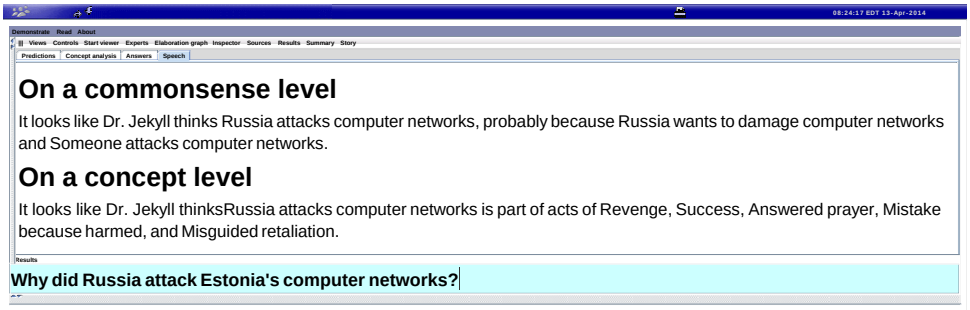

Figure 7: Genesis answers a question about the Russia-Estonia cyber war on several levels.

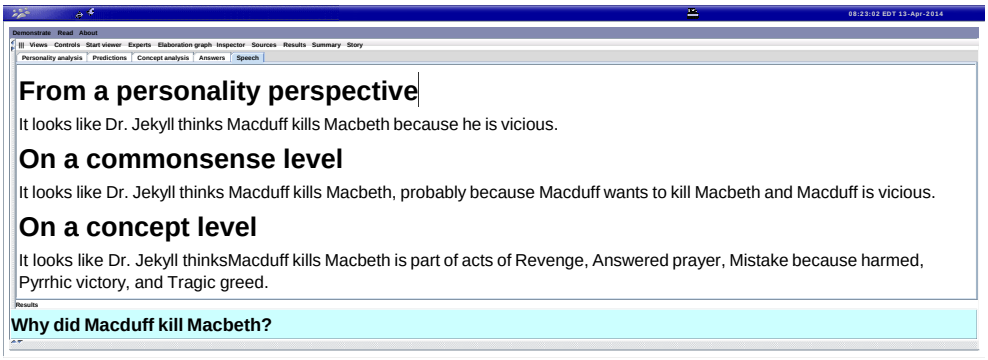

Figure 8: Genesis answers a question about *Macbeth* on several levels including a personality-focused level.

Genesis calculates similarity using concepts

Genesis judges similarity in multiple ways. One way is by using word vectors; another is by using concept vectors. Using concept vectors enables Genesis to see similarities not evident in the words. Two stories may involve, for example, revenge, even though neither uses the word. The comparisons shown in figure 11 are on pairs of short descriptions of conflicts.

The example appears in the MEng thesis of Cary 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 experiment.

As shown in figure 12, a student murders a professor and another student. The Eastern reader has no opinion on why Lu killed Shan until asked if it was because America is individualistic. Then, as shown in figure 13, the Eastern reader, on being asked a question, recalls that he thinks so, which leads to adding that recalled belief to the

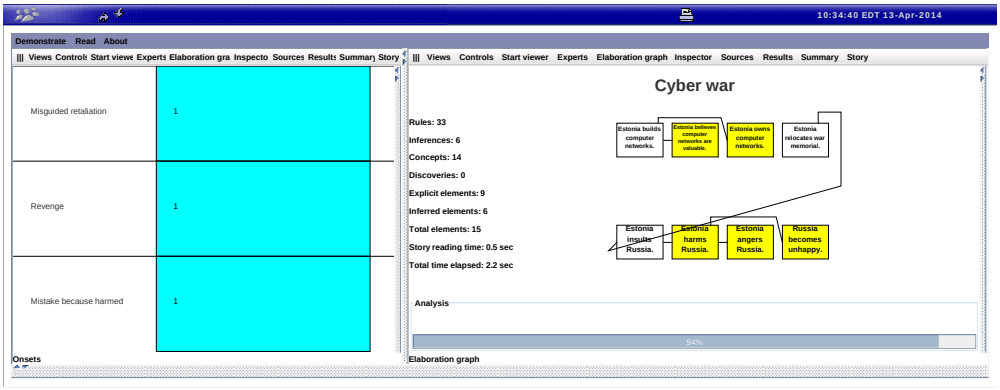

Figure 9: Genesis notes the onset of three possible concepts midstream in Estonia-Russia cyber war.

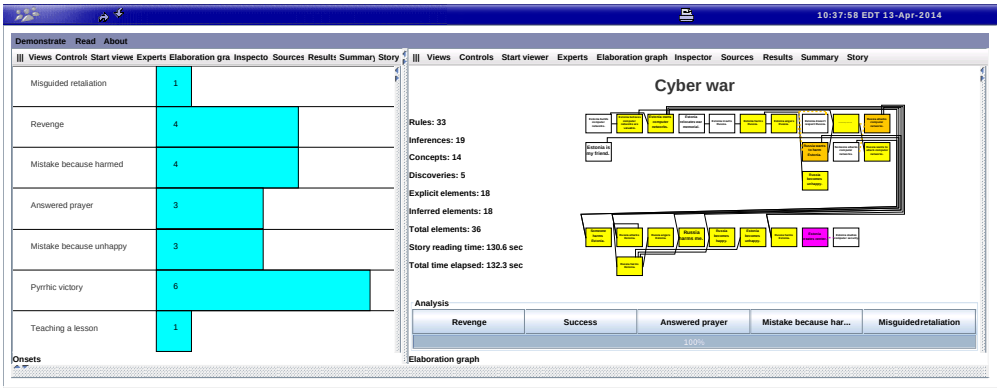

Figure 10: Genesis concludes that all three of the anticipated concepts eventually become realized.

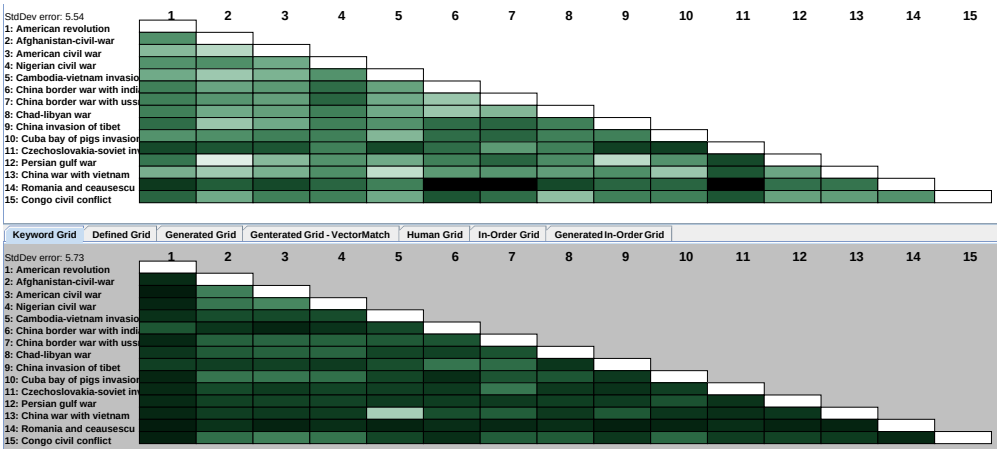

Figure 11: Genesis performs concept-based similarity measurements. Concept-based measurements are shown above and word-based similarity measurements below. White means most similar.

story, when enables connection to the murder. The Western reader has no such belief, so fails to follow the same line of reasoning.

The example appears in the MEng thesis of Hiba Awad (2013).

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 14, Genesis aligns the Tet Offensive from the Vietnam war with the Arab Israeli

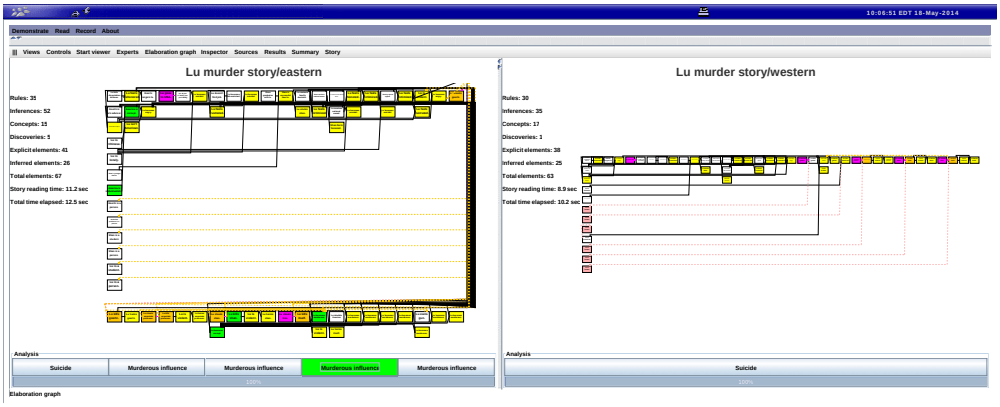

Figure 12: Genesis interprets the Shan murder. The basic interpretation does not connect the murder with America.

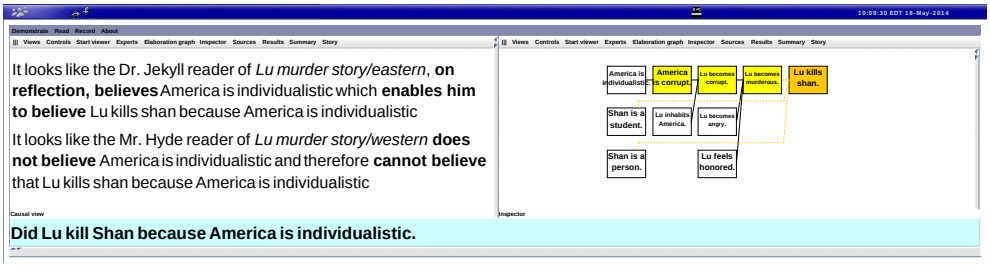

Figure 13: Genesis, stimulated by a question, connects the murder with America's supposed individualistic nature.

war. Such alignment helps see how precedents apply to current events.

Figure 14: Genesis aligns elements in two wars enabling gaps on both sides to be filled.

The example appears in the MEng thesis of Matthew Fay (2012).

Genesis develops summaries

Because Genesis understands stories, Genesis can construct intelligent summaries, such as the one shown in figure 15 for Macbeth, which represents considerable compression relative to the version of the story provided.

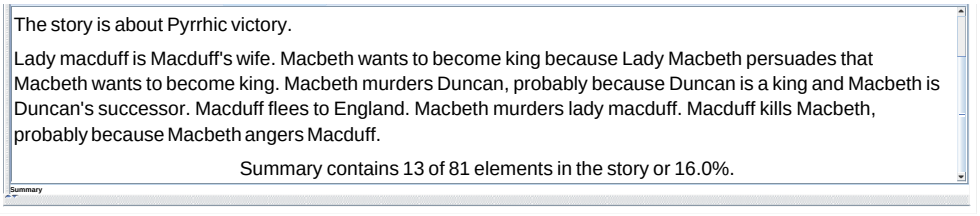

Figure 15: Genesis composes summaries.

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.

In figure 16, 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.

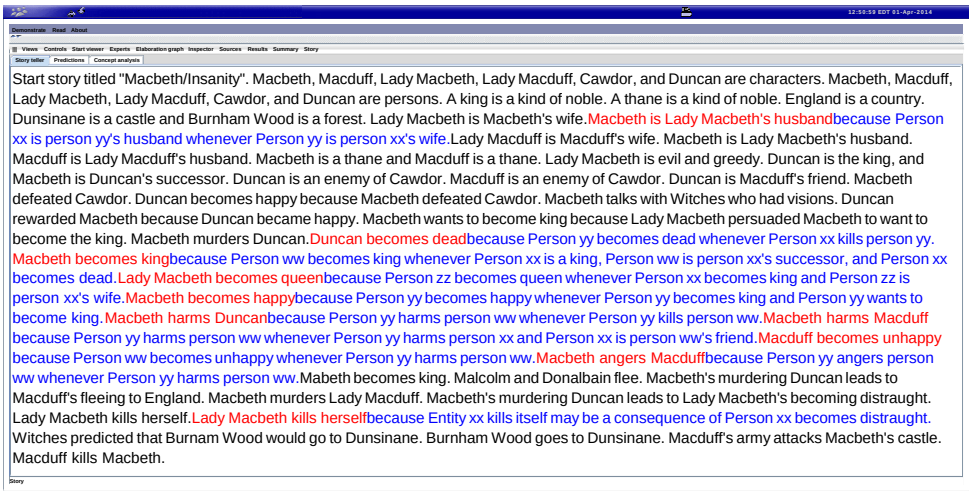

Figure 16: Genesis uses a reader model to determine what and how much to say. Here, Genesis says a lot, because Genesis's model of the reader suggests that the reader does not know much.

Similarly, Genesis can tailor what is said to shape reader opinion. In figure 17, for example, some sentences are emphasized, while others are deleted, so as to make the Woodcutter look good, and everyone else look bad, in Genesis's version of Hansel and Gretel.

The teaching and persuasion examples appear in the MEng thesis of Sila Sayan (2014).

Genesis operates on all of Minsky's six levels

In *The Emotion Machine* (2006), Minsky describes six levels of thinking: instinctive reactions, learned reactions, deliberative thinking, reflective thinking, self-reflective thinking, and self-conscious reflection.

Here are some correspondences between Minsky's levels and Genesis competences:

- Inference rules, the basic rules that produce the basic elaboration graph, are much like instinctive reactions and learned reactions.
- Explanation rules do a kind of deliberative thinking driven by a desire to find ways to produce a more connected elaboration graph.
- Concept patterns examine the elaboration graph and thus perform a kind of reflective thinking.
- Mental models offer some of the capabilities found in Minsky's discussion of self-reflective thinking and self-conscious reflection.

Evidently, Genesis operates on all of Minsky's levels, although not yet performing all the functions described in Minsky's book.

Next steps

Turing concluded his *Computing Machinery and Intelligence* paper (1950) with: "We can only see a short distance ahead, but we can see plenty there that needs to be done." Still true.

Work underway

Genesis is under development by students at various levels, from class projects to PhD theses. The following are representative systems under development:

- A system that composes original stories from precedents.

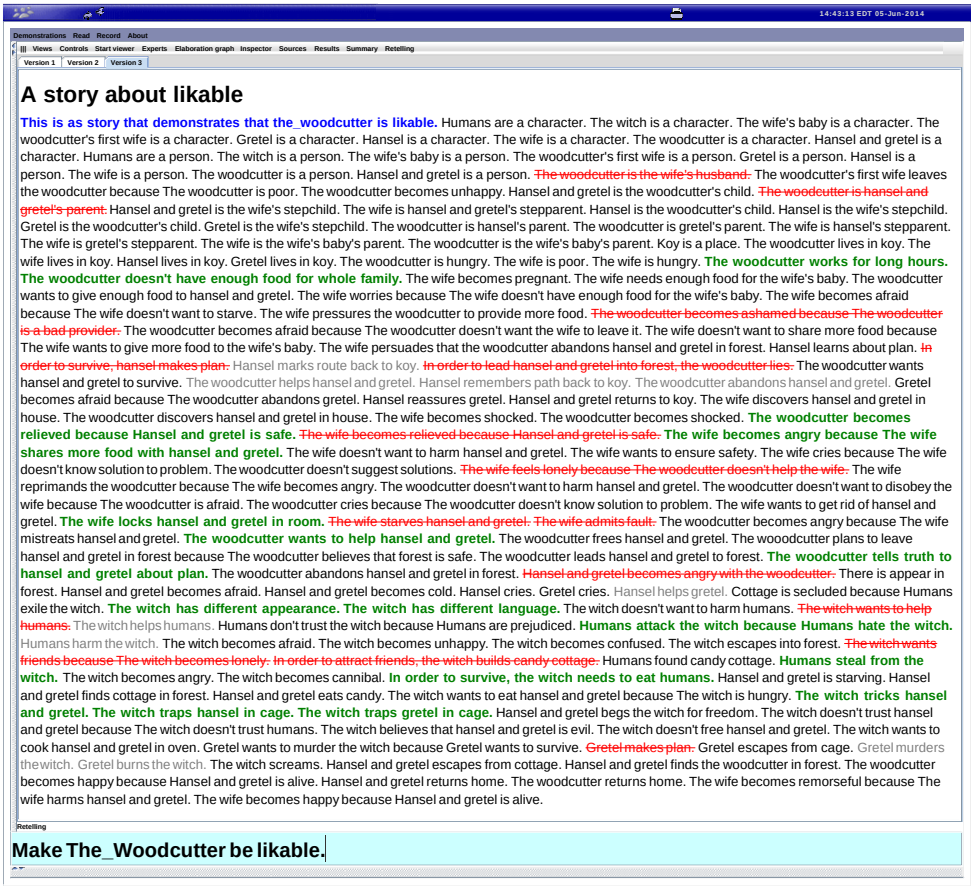

Figure 17: Genesis uses a reader model to determine what and how much to say so as to shape the reader's opinion. Here, Genesis makes one character look good, and emphasizes that goodness, by making the other characters look bad.

- A system that mines literature for actions associated with personality traits.
- A system that tells stories with metaphorical reference to precedents.
- A system that uses story understanding apparatus to plan.
- A system that uses accumulated knowledge to disambiguate verbs.

The CBMM challenge

The central CBMM challenge problem is to understand pictures and video computationally, developmentally, neurobiologically, and socially. Story understanding and telling play a role, of course, because pictures and videos tell stories. One interesting step was taken by Virginia Chiu in an undergraduate research project in 2012. Her focus was on how people describe people. She noted, as expected, that people prefer features that distinguish individuals from other individuals in context, and she noted, as expected, that people prefer features that change slowly, such as physical description, to features that change rapidly, such as location. In a pilot experiment involving 41 descriptions collected from four individuals, she found:

- 31 contained a clothing description, as in *The man with sunglasses wearing a striped blue shirt.*
- 27 contained a physical description, as in *The Asian woman with her hair tied back.*
- 16 contained a location description, as in *The man on the right.*
- 12 contained a manner/action description, as in *The woman who walked past the door to the library.*

Fortunately, from a story telling point of view, Genesis uses Boris Katz's START (1997) to generate English, and START can generate richly descriptive sentences, including sentences with subordinate clauses.

Contributions

What benefits will follow from successful work on story understanding? At least these:

- For CBMM in particular, we will handle stories told in pictures and video.
- For science in general, we will have a better understanding of the key differentiator of our intelligence.
- For applications, our progress on the science side will surely constitute steps toward applications on higher level than we can hope for with today's technology.

References

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

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

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.

Marvin Minsky. Steps toward artificial intelligence. In E. A. Feigenbaum and J. Feldman, editors, *Computers and Thought*. MIT Press, Cambridge, MA, 1961.

Marvin Minsky. *The Emotion Machine*. Simon and Schuster, New York, NY, 2006.

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

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

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.

Alan M. Turing. Computing machinery and intelligence. *Mind*, 59(236):433–460, 1950.

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.

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.