Computer Science and Artificial Intelligence Laboratory

Technical Report

MIT-CSAIL-TR-2009-063 December 17, 2009

Advancing Computational Models of Narrative

Whitman Richards, Mark Alan Finlayson, and
Patrick Henry Winston

massachusetts institute of technology, cambridge, ma 02139 usa — www.csail.mit.edu

Advancing Computational Models of Narrative*

Report of a Workshop held at the Wylie Center, Beverly, MA, Oct 8 – 10 2009

Whitman Richards, Mark A. Finlayson, & Patrick H. Winston
{wrichards, markaf, phw}@mit.edu
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
32 Vassar Street, Cambridge, MA 02139 USA

Introduction

Narratives are perhaps the most effective way to convey to new generations the traditions, knowledge and morals of a culture. They are often used to influence present and future beliefs, as well as cultural norms. One present-day example is the spin placed on reports of the same events by different news outlets — contrast Al Jazeera and the New York Times — that serve to inflame opposing passions of the Arab and Western world. Consider also the history textbooks given to Israeli and Palestinian children, which differ significantly in their accounts of those peoples (Bar-On & Adwan 2006). Analyzing these phenomena by building formal models is an important step toward understanding how members of a culture might reason about the impact of internal and external events and actions, and therefore eventually mitigating and controlling the negative effects. Surprisingly, given the 50-year history of Artificial Intelligence, there have been few research programs that have investigated narrative itself from a computational viewpoint. This first

* Sponsored by the AFOSR under MIT-MURI contract #FA9550-05-1-0321. Special thanks to Maria Rebelo for administrative support.

workshop on computational approaches to understanding narrative had three main objectives:

- Evaluate the "state-of the art" in modeling narrative
- Explore the scope and dimensions of narrative itself
- Begin to build a community focused on computational narrative

To address these issues, an interdisciplinary group was assembled. Included were computer scientists, psychologists, linguists, media developers, philosophers, and story tellers (See Table 1 below.) Abstracts showing each individual's perspectives were circulated prior to the workshop, and updated versions are included in Appendix 2.

Presentations were organized into four groups, each followed by a panel discussion (See Appendix 1). On the morning after these presentations there was an extended discussion of issues related to narrative understanding and how computational approaches might foster insights and facilitate analyses. The sections that follow recap the main issues, potential advances, lacunae, and next steps to build a larger community.

Nicholas Asher, *University of Texas*	Henry Lieberman, *MIT*
Neil Cohn, *Tufts*	Terrence Lyons, *AFOSR*
Michael Cox, *DARPA*	Erik Mueller, *IBM*
Mark Finlayson, *MIT*	Srini Narayanan, *ICSI and UC Berkeley*
Ken Forbus, *Northwestern University*	Whitman Richards, *MIT*
Pablo Gervás, *Univ. Complutense de Madrid*	VS Subrahmanian, *University of Maryland*
Jerry Hobbs, *USC ISI*	Reid Swanson, *USC ICT*
Ian Horswill, *Northwestern University*	Emmet Tomai, *University of Texas Pan American*
Ray Jackendoff, *Tufts*	Bart Verheij, *University of Groningen*
Jack Jackson, *Naval Postgraduate School*	Patrick Winston, *MIT*
Jay Keyser, *MIT*	Michael Young, *North Carolina State University*

Table 1: Workshop Participants

Representations for Narrative

Computational modeling requires a well-defined statement of the problem (or problems) to be solved. An obvious first step, therefore, is to agree on what defines the character of *narrative*, or at the very least, to identify its boundaries. Classically, narrative is usually cast as a succession of happenings within a setting (or context), and a plot involving a set of characters. However, there are many different representations that could be invoked, especially if the narratives have quite different goals and objectives. These differences became quickly obvious:

most workshop participants came with a representation in mind that was tailored to their own specific research interests (even if it was only implicit). Despite these differences, however, there were some common elements.

Consider Verheij's representation of a narrative shown in Figure 1. This figure embodies what might be considered the kernel of all the narrative representations considered. At the top, the large arrow represents the main progression of the story: a linear set of events that proceeds, one after another, from start to finish. Beneath that arrow are smaller arrows, representing smaller portions of the story that could be considered as stories in their own right. This nesting can proceed quite far, until, finally, the stories ground out in some common-sense knowledge of the world, suggesting connections between argumentative and narrative elements (Bex et al. 2007, Bex, Prakken & Verheij 2007).

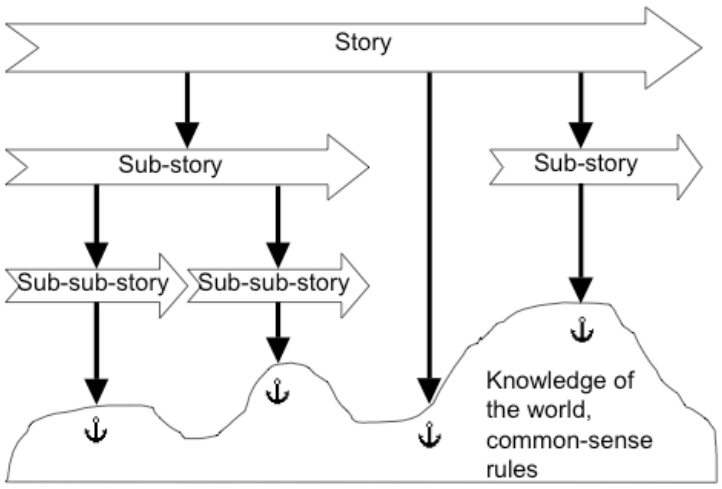

Figure 1: Verheij's visualization of hierarchical narrative grounded in common-sense, adapted from the representation of Wagenaar, Van Koppen and Crombag (1993)

Hence, to first order, there are three common denominators amongst representations considered: (1) narratives have to do with *sequences of events,* (2) narratives have *hierarchical structure,* and (3) they are (eventually) grounded in a *commonsense* knowledge of the world.

That the event sequences may be in some sense *non-linear* is illustrated by Gervás's depiction, shown in Figure 2. There, consider the short pieces of text on the left to be abstract "discourse segments" used to construct a narrative. On the right, the relationship between actors for each action segment are displayed progressively at each major time step.

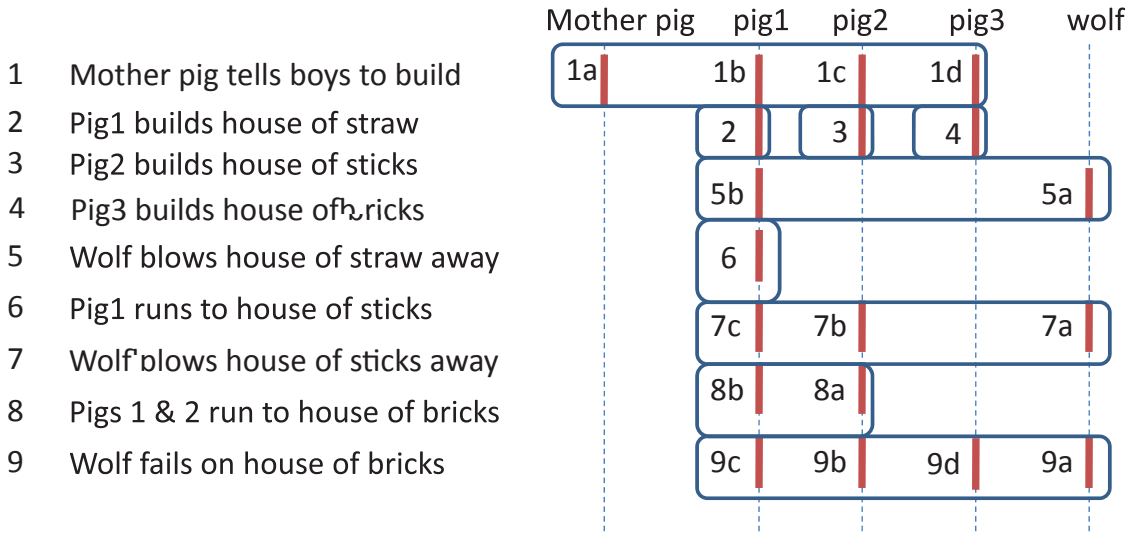

Figure 2: Gervás's representation of the multiple points of view of narratives

The representation on the right is one way of capturing the major event structure of a narrative, and is especially important when considering that the same story may be told from *multiple points of view*. This points out an interesting computational problem, namely, to reconstruct the structure shown on the right from the discourse segments shown on the left.

With these major characteristics of narrative in mind, we emphasize that all are independent of lexically-based language. Cohn made this point forcibly with his examination of the narrative structures within a form that can be purely visual: the visual language used in comics. This visual language conveys complex sequences of events, can have hierarchical constituency structure (as illustrated in Figure 3), structural ambiguities, head/modifier relationships, and involve multiple characters and points of view. This visual language can tell a story without ever using a single word.

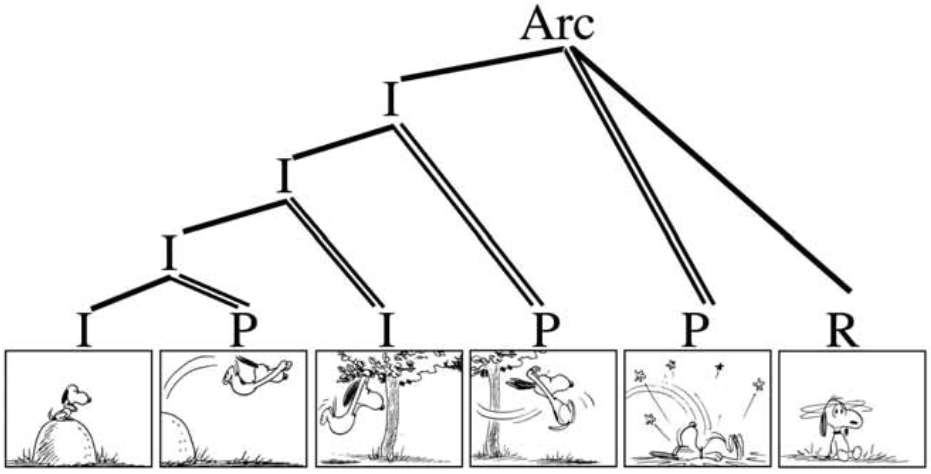

Figure 3: Illustration of hierarchical structure within a purely visual medium (Cohn 2009)

From these fairly non-controversial starting points researchers expanded in multiple directions. For example, Verheij elaborated Figure 1 to provide a scheme for modeling legal narratives. This requires an analysis of the details of narrative: how evidence is assembled and presented before a judge and jury. To convince a jury to convict a defendant, a prosecutor must deliver a well-formed narrative, with the assembled components nailed down, and with reasons for the various acts made compelling. Verheij illustrated how a Wigmore chart (show in Figure 4) can be used to analyze which parts of a narrative support, explain, or contradict – in other words, how the different parts *fit together*.

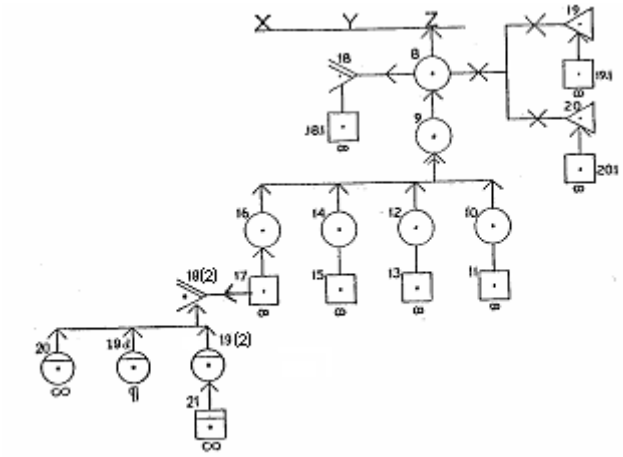

Figure 4: A chart (from Wigmore 1931) that shows how different parts of the prosecution's narrative (for a particular charge of murder) support each other

Taking yet another direction, Mueller's aim was to build machines for narrative understanding, as verified by questions answering (Mueller 2003, 2004, 2007). To do this, his process, illustrated in the flowchart in Figure 5, relies on a detailed, axiom-based *commonsense knowledge base* coupled with a logic theorem prover. Mueller's approach emphasizes that story understanding builds on implicit as well as explicit knowledge possessed by the listener. In the Three Little Pigs story, for example, one knows that a house of straw is fragile compared with one built of bricks, and a wolf is already assumed to be predatory. Clearly, if a story is to be "understood" then the implications of the events and their sequences should be recoverable, even if these implications are not explicitly mentioned.

Hobbs also tackles the story at this level of understanding, casting story understanding not as theorem proving, but as *abductive inference* to a plan that best explains the story (Hobbs 1990, Chapter 2). Hobbs's view is that the characters in the story are cognitive agents engaged in planning mechanisms; namely, that they are trying to achieve goals in the face of some obstacles, and

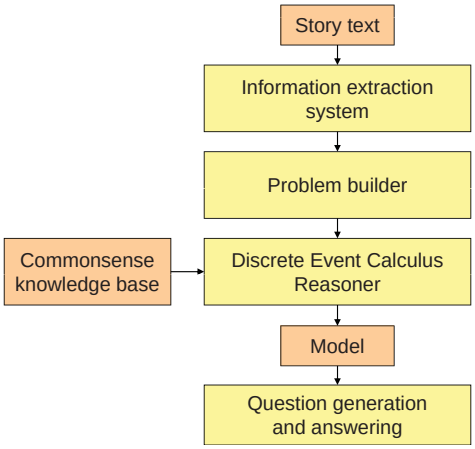

Figure 5: Flowchart of Mueller's story understanding process

proceed to break the main goal into subgoals, and those subgoals into further subgoals, and so forth, until a sequence of achievable actions leads to the desired outcome. The agents operate in an uncertain, partially-obscured world, as we do, and so must monitor the state of the world and adjust to keep their plans on track.

Young's approach on narrative generation is related (see Figure 6). The problem is adapting a pre-constructed narrative (for, say, a computer game) to a player's unpredictable actions (Young 2007). Interestingly, analysis at such a detailed level has opened up an avenue toward customizing aspects of the "suspense" of the narrative. If one takes into account that narrative follows natural Gricean conditions, like any communicative text, then you can reveal or conceal information to increase the suspense.

A few discussed work on moving toward more abstract, generalized representations, as opposed to a more refined, detailed analysis. Forbus and Tomai's approach to reasoning with moral narratives used, for one branch of their reasoning system, analogical transfer and generalization. These processes take into account the higher-level structural similarities between two stories, and try to transfer consequents from one to the other causal structure. Narayanan and Lieberman also stressed the role of common sense. Narayanan discussed the role of physical knowledge and frame semantics (see Figure 11 and Loenneker-Rodman & Narayanan 2010, Feldman & Narayanan 2004, Feldman 2006) whereas Lieberman was relying on OpenMind (Kim, Picard, & Lieberman 2008).

Finlayson's work expands on the constraints imposed by the higher-level structure. The approach is to derive the high-level structure by comparing, contrasting, and merging a large set of stories (Finlayson 2009). The structures that result are Propp-like morphologies that encode a finite state grammar – an example is shown in Figure 7. Note that one use of this representation is to make vivid the higher-

level analogies, specifically the similarities and differences between stories – an important component of cultural reasoning.

Despite the relative few working on the "higher-level" representation of narrative, all agreed that this was an extremely important area to concentrate on. Mueller noted that the first thing he would work on, if the complexity of reasoning over his knowledge bases were overcome, would be to get at the nature of the story more,

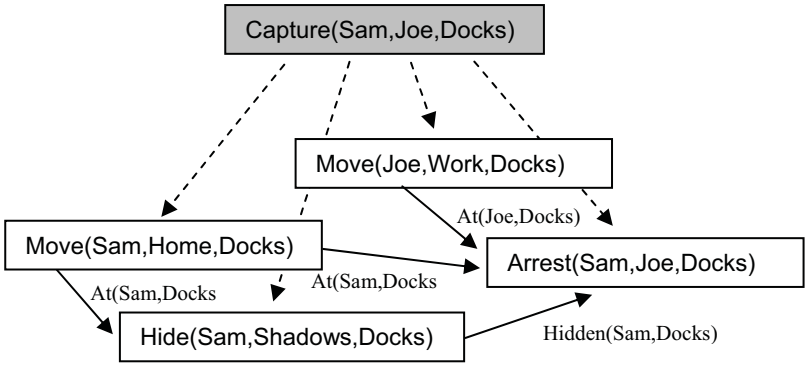

Figure 6: Example of a plan fragment, meant to be a part of a larger story-plan structure (Young 2007)

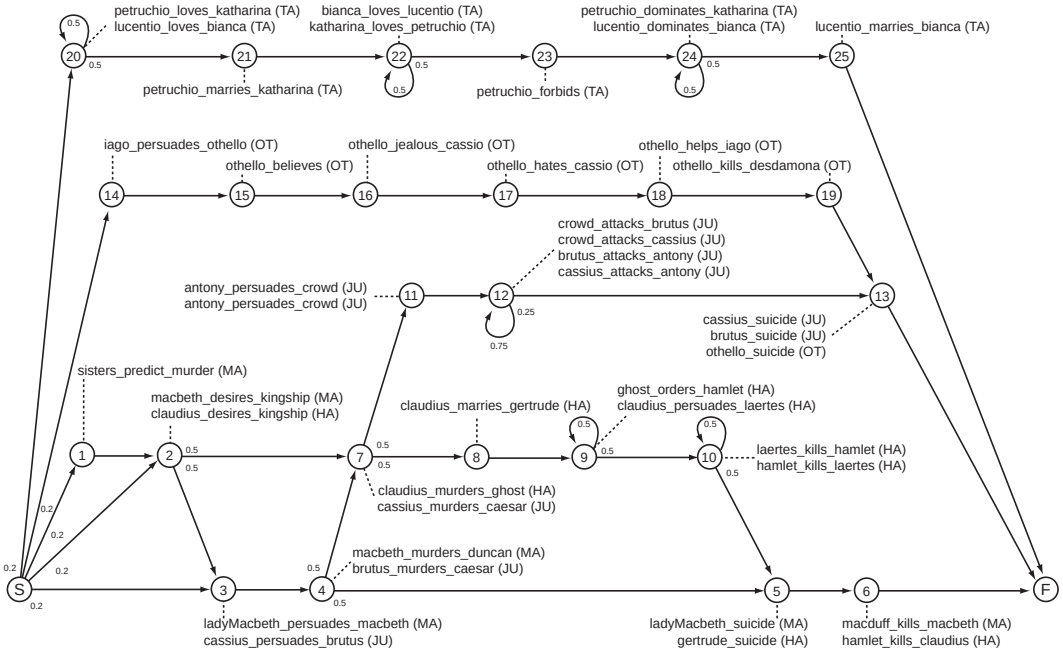

Figure 7: An example narrative morphology, illustrating a level of description higher than an individual story (Finlayson 2009)

and explore representations by Shank, Dyer, Lehnert and others that treat the story at the more abstract level. Asher was emphatic that isolating and understanding the higher-level structure was one of the primary motivations for doing low-level annotation.

As can be seen, most of the representations agreed on a few fundamental points. First, event structure was generally accorded a central place. Second, most every representation had the facility to represent hierarchical structures in some form. Third, almost every representation or model included, either implicitly or explicitly, commonsense knowledge. From this common ground there was divergence. Some focused on the microstructure of narratives, e.g., the axiomatic commonsense knowledge, planning structures, argument structures, or physical commonsense knowledge. Others focused on macrostructure, such as at the causal level, a generalized plot unit level, or at the level of unified meaning.

In sum, the principal obstacles to defining narrative in a manner suitable for computational studies – or toward a theory of narrative – came not from a disagreement about core structures at play. Rather, the differences came in the goals of the narrative, how these different goals impacted the representation, and whether, and in what way, micro- and macro-levels should be represented and used.

Dimensions of Narrative

The second major issue addressed by the workshop participants was the nature of narrative, its scope, and which types of narratives warranted study. Table 2 lists types of narratives mentioned during the workshop, split into somewhat arbitrary categories.

Formal	Somewhat Formal	Traditional	Informal / Spontaneous
Novels	Some poetry	Folktales	Anecdotes
Movies	Blogs	Fairy Tales	Diaries
Biography	Computer games	Fables	Gossip
Plays	Comics	Fabula	Some jokes
Case studies	News	Myths	Water-cooler talk
Legal Argument	Sitcoms	Legends	Urban legends

Table 2: Selection of types of narratives mentioned at the workshop

The purpose of the table is not to lay out a precise ontology of narrative, but rather to illustrate the range of phenomena that should be considered. There are those artifacts that sprung to everyone's mind immediately on the word *narrative*: objects that fall into more formal categories, like novels, plays, or news. But, as Horswill pointed out, it is helpful to keep in mind other sorts of exemplars of a less formal, and more spontaneous sort, such as diaries, gossip, or water-cooler talk. These are also narratives, and perhaps form the main bulk of people's experience with narratives from day-to-day.

Clearly, with such a range of genre, narratives have many dimensions. First and foremost is the *event structure*, namely, narratives are about something happening.

Second, there was the *narrative structure*. Most participants thought event structure was distinct and separable from narrative structure. But the group could not achieve consensus on what structures made something a narrative. For example, Hobbs presented *chronicles*, which are rolls of unrelated events, and said they were perfectly well-formed lists, but not narratives. To counter this, Keyser brought up the example of picaresque novels, which are just a series of unrelated adventures of a single character, and noted that most sitcoms are just reduced forms of the picaresque novel. Jackendoff, to second the observation, noted that in music, variation movements don't have much of a broad "narrative arc" like other forms, but are a musical form of the picaresque novel.

The next most important dimension, from the point of view of the discussion, was the *purpose* of a narrative. Narratives can be used for many things, for entertainment, for argumentation, for political propaganda (e.g., Lakoff 2008), etc. It was noted by Asher that often the purpose of the narrative is orthogonal to the structure of the narrative itself. To illustrate this point, he contrasted a novel, which might use indirection, misinformation and surprise to achieve its goals, with a legal argument (as discussed by Verheij) that should be straightforward, matter-of-fact, common-sensical and lacking in surprise to achieve its goal of overcoming a reasonable doubt. Consider also the goal of conveying moral issues via a story, or through mythical analogy versus the goal of telling a joke, or perhaps more extreme, a scientific account of the evolution of homo-sapiens, or of a "brain". In each of these examples, the knowledge and reasoning abilities required differ considerably. Related to the purpose was the idea of the function of narrative in different societies: Horswill proposed that the main function of narrative was to mediate individual and group identity. Jackendoff countered that, however, if you look across the whole range of narrative examples, you find many different functions.

Fourthly, the role of the *listener* must be addressed. As noted by Cox, the narrator must be tuned to the abilities and interests of the listener if he wants to achieve his purpose. How the story is presented (i.e., with explanatory remarks of footnotes, etc.) may need to be revised depending upon who is listening. Additional media might be introduced – for example, handouts in a lecture versus an unadorned reading of a poem.

Related to the above, Horswill pointed out a dimension that was little treated at the workshop, the issue of the emotional *impact* of the narrative on the listener. Subrahmanian stressed this dimension as well, noting it was especially important in light of the varied purposes to which one might apply a narrative. He gave the example of trying to write a book today on a well-tread subject, say, for example, the Rwandan genocide. To make the book "worth reading by anybody", it would need to say a lot more than just particular things occurred at particular times, be more than just a list of events. It would need to argue, hypothesize, speculate, surprise, suggest; it would have to deal not only with events as they happened but possible events and states of the world; it would have to be the right length so as to not bore the listener, but not leave them dissatisfied.

Also lightly tread, but still mentioned, was the *medium* of the narrative. While there was general consensus with Cohn's view that narratives are not tied to lexical forms, it was often noted that a change in the medium results in a change in the narrative. The workshop participants felt that these "details" and issues were critical, but had a very difficult time articulating how exactly this particular dimension should inform computational models.

Finally, Keyser several times emphasized the importance of the meaning of the narrative – the *narrative arc*. No representation specifically addressed this point, namely what a story *means* – its moral or overall message. In a related point, Asher raised the issue of how the nature of the message might recast the narrative framework.

Hence the general consensus was that computational studies should explore a range of media, with narratives of different genre and goals. To this end, there was a strong recommendation that the present computationally-oriented community present at the workshop should be expanded (see below.)

Fundamental Problems

In the preceding sections, three main problems were addressed: why narrative?, what constitutes narrative?, and what are the appropriate representations? While the discussion in the report sketched potential answers to these questions, many issues remain.

Why narrative? There was substantial discussion of what makes narrative important and worthy of study. As Horswill and Jackendoff pointed out, participants were unable to give a concise and cogent reason why narrative *itself* should be an object of study, rather than a variety of other cognitive or social processes. Each participant had their own specific examples of applications and interests, but there was no over-arching, knock-down reason. Participants generally agreed with Cohn's assertion that narrative is a ubiquitous way to structure, package or order information such that we can understand it, while disagreeing as to whether narrative itself was an epiphemenon of some deeper process, say, analogy (Forbus), planning (Hobbs), or social interaction (Horswill). Tomai noted that narrative seems to be a form especially well suited for communicating complex sequences of events over a low-bandwidth channel, by taking into careful account different constraints of language and context. But all of these observations still beg the question of what is special about narrative *in particular* – there are many ways of structuring information without it being a narrative.

What exactly is narrative? Hand in hand with the previous problem, although numerous examples and dimensions were identified, no one was able to truly define what a narrative is, to give a procedure for distinguishing good narratives from bad, or distinguish narratives from non-narratives. Keyser asserted that a good narrative is one which is "about something", something that can be extremely simply stated. Hobbs and Gervás countered that often a critical view of literature holds that the more interpretations a text has, the better. Jackson recounted that this was the explicit technique taken by the script writers for the recent popular television show *Battlestar Galactica*: that they tried to keep it as open-ended as possible. Verheij noted that in legal psychology it has been suggested that consistency, non-ambiguity, and chronological ordering make a story better in the sense that stories with these properties are more quickly believed to be true.

What are the appropriate representations for narrative? Here too we are left in an incomplete state. It is clear from the survey that no presented representation or system spanned the whole range of narrative levels identified, and there was no attempt or even a sketch of a unified representation or representational architecture. Cohn notes that this lack of unified models may be because the answer to the question of representation is, at least in part, conditioned on the answer to the previous question, namely, our understanding of what narrative is.

Nevertheless some first steps were taken to address these issues: the MoralDM system (Dehghani et al. 2008), presented during Forbus's talk, integrates representations at levels both above and below the generally-accepted event structure middle level. For the level below, they used first-principles reasoning

(bottom branch in Figure 8), and for the level above, causal reasoning by analogy (top branch).

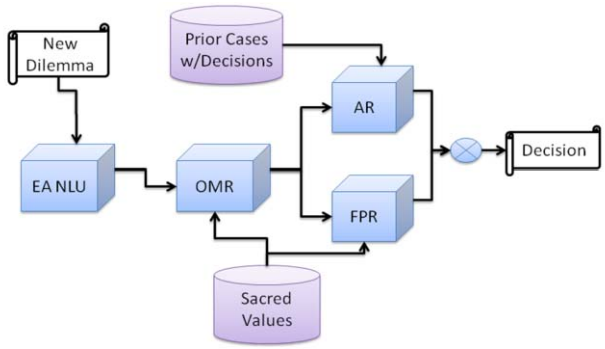

Figure 8: The MoralDM architecture, reasoning at both the detailed (lower branch) and abstract (upper branch) levels

Facets of understanding narrative: Putting aside both the representation and narrative definition issues, an extended discussion of what would be required for full understanding of narrative raised the following sets of questions:

(i) What makes narratives different from a list of events, or facts? Jackendoff hammered on this point throughout the conference: what is it about the structure of narrative that makes it cohere? Asher asked too, what is special about the arrangement of the items in the discourse that makes it a *narrative*, rather than something else?

(ii) Forbus asked if understanding narrative first requires us to understand common-sense reasoning? Verheij asked about the effect of small changes in the story for its use and interpretation, noting that logic-based AI is traditionally quite sensitive to initial conditions.

(iii) How are stories indexed and retrieved? Is there a "universal" scheme for encoding episodes? (Jackendoff made here an analogy to the idea in language of Chomsky's universal grammar)

(iv) Jackendoff asked if the composition of the narrative is more dependent on its purpose, or rather on the understanding of the underlying event structure? The purpose, function, and genre of a narrative seems to have an enormous impact on the form and content of the story – what are those impacts? Does the genre of the story change not only the surface form, but the underlying representations used?

(v) Asher and Keyser in particular asked how moral decisions and content relate to narrative structure. Are their systematic differences in the formal properties of narratives from different cultures?

(vi) Keyser asked what is the number of narrative arcs? What are the possible story lines? Is conflict almost always in ingredient? Is there a recipe, such as Campbell's Hero's Journey, or a Proppian morphology such as derived by Finlayson? An argument against this view was raised – namely if the listener's stance must be included in the design of a story, then the many factors come into play, including genre and generation, etc. At an even higher level of abstraction, Keyser pointed out that no representation discussed or implied at the workshop dealt with what a story means, in the sense that War and Peace is about "people's inability to control events," or Don Quixote's true message is that "to fully enjoy life you have to be a little crazy."

(vii) What does the microstructure and macrostructure of a narrative look like? Are these representations unique? How is narrative structure different from discourse or rhetorical structure?

(viii) Where does the narrative technique of *characterization* fit in these schemes? Tomai noted that there was much discussion of two of the classical parts of narrative, namely, plot and setting, but very little about the third, *character*.

(ix) What are the representations that underlie the extraction of schema from the blooming, buzzing confusion of the world? Induction from raw experience seems difficult, and Narayanan speculated that some of the mechanisms are built-in and evolutionarily quite old, e.g., the idea of 'starting' and 'stopping' as a bootstrapping mechanism for extracting events from experience. What are the cognitive principles that permit listeners to assign schematic structures to heard narratives? What aspects of these structures and principles are specific to narrative, and what aspects come from general-purpose cognition (or elsewhere)? Winston asked if the full comprehension of a narrative requires internal pictorial images.

(x) How is the listener convinced that a story is true? A common case is when one is trying to evaluate different, conflicting accounts of real events presented by news media. It is clear that the current computational work in legal arguments and narrative is relevant, where testimonies often conflict. If a story (or argument) is compelling, to what extent (and why) might false accounts be perceived as truthful? Lyons was particularly concerned about narratives with respect to the confirmation bias – he noted this is a concern not only for legal argument, but for medical reports or strategic reasoning.

(xi) How do stories help drive decision-making? Analogy is the clearest computational route, but are there others? Lieberman asked how one reasons by analogy with conflicting precedents, and Gervás asked what one does when different aspects of structure within the same story conflict. Jackson noted that if you have multiple analogies, when you decide on a course of action (because of other information), you can justify your choice with the corresponding analogy. Verheij's work in particular bears on how narratives affect decisions, since the narratives the defense and prosecution tell directly influence the judgment they issue.

(xii) What insights do narratives give about causality? How do stories help us deal with or structure our world? How do we extract causal models from narratives?

(xiii) How do we evaluate computational models and themes of narrative? Questions and answers? Ability to generate compelling stories? If evaluation is to be against a databank of stories, then what should be the normal form for stories? Just how will restrictions on annotation affect the ability to find analogies, similarities, etc?

(xiv) Narrative structure aside, what makes a "good" story? (Keyser) How does the judgment of "good" depend on the genre or purpose of the story? Why do we like stories?

Steps for the Future

Two initiatives were applauded by all: (1) a second workshop in one year and (2) a story data base, analogous to the Penn Treebank used for evaluating statistical parsers.

Second Workshop: Unanimously, the participants agreed that the workshop was a boost to understanding narrative, by bringing together a variety of approaches, showing links and differences. It was felt that the community was fragmented, and needed to be encouraged and grown. Many of the participants had not previously met, and consequently a variety of perspectives and approaches were new to large segments of this small group. Hence the workshop was an important first step to creating a community of researchers studying narrative from a computational viewpoint. A second workshop would be the obvious net step toward establishing a larger, still broader community. We visualize the next workshop as doubling in size, and including several areas not represented, such as game-theory approaches to narrative, studies of gossip and rumor, and especially narrative theory researchers from the humanities. More thought is also needed to reach an

agreement on methods for evaluating story understanding, as well as various experimental paradigms. In addition to broadening the scope of participants, a second workshop is needed to investigate whether a new community should be set up (with its own annual meetings and publication vehicles), or if, as Winston suggested, the participants are naturally a subset of an already established community. In particular, it should be investigated whether the participants feel a publication venue directly associated with the area would be appropriate.

Catalogue of Problems and Applications: Richards suggested that a list of potential applications of narrative – big problems on which narrative might give one traction – would be of great use to motivating work and securing funding. While this effort may not have to have its own separate committee, it was thought that some attention should be paid to assembling such a list. A few candidates for inclusion were offered:

1. Using narrative to filter incoming information, to interpret it, impose a structure on it, all sensitive to your goals (Cox)
2. Detection and production of propaganda (Winston)
3. Understanding and influencing other cultures (Richards, Jackson)
4. Helping others tell their own stories, say, for entertainment purposes (your summer vacation) or institutional knowledge capture (Lieberman, Cox)

Data Bank: The majority of participants were engaged in computational approaches to text based story understanding. At present, there is no shared corpora of stories – a necessary tool if one is to compare successes and strengths of various approaches. Hence it was proposed to create a story databank.

One important property of the story databank would be to provide at least one translation of text into an agreed-upon format, such as Finlayson's set of semantic annotations or Asher's discourse representations. It was generally agreed that this annotation is an extremely time consuming and delicate process, no matter if done manually (e.g., Cohn, Mueller, Asher), automatically (Forbus and Tomai), or semi-automatically (Finlayson). Gervás noted that the assembly of a corpus had been discussed quite a bit in the story generation field, but beyond agreeing that a corpus would be useful, there was no agreement on what exactly to put in the corpus or how it was to be represented. Hobbs expressed grave concerns that to get high inter-annotator agreement any representations for such a corpus would have to be too simple to be helpful.

The general consensus, though, was that, despite the difficulties, there should be some attempt to make a story databank. Fortunately, the Story Workbench (Finlayson 2008), ideally in concert with Tomai's EA NLU system (Tomai &

Forbus 2009), will shortly be ready to provide such a translation of stories into a form more suitable for computational analysis modeling. Narayanan also mentioned that FrameNet from Berkeley might be of use, as it already has small numbers of short, frame-annotated stories

The story databank will not only need the stories themselves and their formal representations, but formal representations of evaluation metrics. One type of metric that is immediately accessible are the question answering metrics used by Mueller, and it was suggested that the story databank include lists of questions (with answer keys) that a story understanding system should be able to answer. But more thought needs to be put into different types of metrics that might be used.

Another element that possibly could be included in the databank would be a heterogeneous set of resources that can be used to *generate* narratives, potentially in multiple different media.

This effort will be led by Gervás, with a tentative committee that includes at the very least, Mueller, Verheij and Finlayson.

Bibliography

Asher, N. (1993). *Reference to Abstract Objects in Discourse.* Dordrecht, Kluwer Academic Press.

Asher, N. and A. Lascarides (2003). *Logics of Conversation*. Cambridge, Cambridge University Press.

Bar-On, D. and S. Adwan (2006). The Psychology of Better Dialogue between Two Separate but Interdependent Narratives. *Israeli and Palestinian Narratives of Conflict: History's Double Helix*. R. I. Rotberg. Bloomington, Indiana University Press**:** 205-224.

Bex, F. J., Braak, S. W. v. d., Oostendrop, H. v., Prakken, H., Verheij, B., & Vreeswijk, G. (2007). "Sense-making software for crime investigation: how to combine stories and arguments?" *Law, Probability, and Risk*, **6**, 145-168.

Bex, F. J., Prakken, H., & Verheij, B. (2007). Formalizing argumentative story-based analysis of evidence, *Proceedings of the Eleventh International Conference on Artificial Intelligence and Law* (pp. 1-10). New York: ACM Press.

Chwe, M. S.-Y. (2009). "Rational Choice and the Humanities: Excerpts and Folktales." *Occasion* **1**(1): http://occasion.stanford.edu/node/9.

Cohn, N. (2009). Foundations of Visual Language Grammar.

Cox, M. and A. Ram (1999). On the intersection of story understanding and learning. *Understanding language understanding : computational models of reading*. A. Ram and K. Moorman. Cambridge, MA, MIT Press / Bradford Books**:** 397-434.

Dehghani, M., E. Tomai, K.D. Forbus, and M. Klenk (2008). An Integrated Reasoning Approach to Moral Decision-Making. *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence*. (held in Chicago, IL), Menlo Park, CA, AAAI Press**:** 1280-1286.

Fayzullin, M., V. S. Subrahmanian, M. Albanese, C. Cesarano, and An. Picariello. (2007). "Story Creation from Heterogeneous Data Sources." *Multimedia Tools and Applications* **33**: 351-377.

Feldman, J. (2006). *From Molecule to Metaphor*. Cambridge, MA, MIT Press.

Feldman, J. and S. Narayanan (2004). "Embodied meaning in a neural theory of language." *Brain and Language* **89**: 385-392.

Finlayson, M. A. (2008). Collecting Semantics in the Wild: The Story Workbench. *Naturally Inspired Artificial Intelligence* (Technical Report FS-08-06, Papers from

the AAAI Fall Symposium). (held in Arlington, VA), Menlo Park, CA, AAAI Press.

Finlayson, M. A. (2009). Deriving Narrative Morphologies via Analogical Story Merging. *New Frontiers in Analogy Research* (Proceedings of the Second International Conference on Analogy). (held in Sofia, Bulgaria), Sofia, New Bulgarian University Press**:** 127-136.

Gervás, P. (2009). "Computational Approaches to Storytelling and Creativity." *AI Magazine* **30**: 63-70.

Hobbs, J. R. (1990). Literature and Cognition. *CSLI Lecture Notes*, Number 21. Center for the Study of Language and Information.

Horswill, I. (2007). Psychopathology, Narrative, and Cognitive Architecture. *Intelligent Narrative Technologies* (Technical Report FS-07-05, Papers from the AAAI Fall Symposium). (held in Arlington, VA), Menlo Park, CA, AAAI Press.

Jackson, L. A. (2009). Narrative Paradigm in Cultural Geography Modeling. *Proceedings of the Human Behavior-Computational Modeling and Interoperability Conference*. (held at Oak Ridge National Laboratory).

Kim, K., R. W. Picard, and H. Lieberman (2008). Common Sense Assistant for Writing Stories that Teach Social Skills. *Proceedings of the Twenty-Sixth ACM Conference on Computers and Human Interaction*. (held in Florence, Italy), New York, ACM Press**:** 2805-2810.

Lakoff, G. (2008). *The Political Mind*. New York, Viking Press.

Lascarides, A. and N. Asher (2007). Segmented Discourse Representation Theory: Dynamic Semantics with Discourse Structure. *Computing Meaning*. H. Bunt and R. Muskens. Amsterdam, Kluwer Academic Publishers. **3:** 87-124.

Loenneker-Rodman, B. and S. Narayanan (2010). Computational Approaches to Figurative Language. *Cambridge Encyclopedia of Psycholinguistics*. Cambridge, Cambridge University Press.

Mueller, E. T. (2003). Story understanding through multi-representational model construction. *Proceedings of the HLT-NAACL Workshop on Text Meaning*. G. Hirst and S. Nirenburg. East Stroudsburg, PA, ACL Press: 46-53.

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

Mueller, E. T. (2007). "Modeling space and time in narratives about restaurants." Literary and Linguistic Computing **22**: 67-84.

Swanson, R. and A. S. Gordon (2008). Say anything: A massively collaborative open domain story writing companion. *Proceedings of the First International Conference on Interactive Digital Storytelling* (Lecture Notes in Computer Science, Vol. 5334). (held in Erfurt, Germany), Berlin, Springer**:** 32-40.

Tomai, E. and K. D. Forbus (2009). EA NLU: Practical Language Understanding for Cognitive Modeling. *Proceedings of the 22nd International Florida Artificial Intelligence Research Society Conference*. (held at Sanibel Island, FL), Menlo Park, CA, AAAI Press.

Verheij, B. (2005). *Virtual Arguments: On the Design of Argument Assistants for Lawyers and Other Arguers*. The Hague, The Netherlands, TMC Asser Press.

Wagenaar, W. H., P. J. van Koppen, and H.F.M. Crombag (1993). Anchored Narratives: The Psychology of Criminal Evidence. New York, St. Martin's Press.

Wigmore, J. H. (1931). *The Principles of Judicial Proof*. 2nd Ed., Boston, Little, Brown and Company.

Young, R. M. (2007). "Story and discourse: A bipartite model of narrative generation in virtual worlds." *Interaction Studies* **8**: 177-208.

Appendix 1: Meeting Agenda

Thursday, October 8

5:00-7:00	Cocktails	
7:00	Dinner	

Friday, October 9

8:30 – 9:00	Opening Remarks, **Whitman Richards** & **Patrick Winst**on, *MIT*	
9:00 – 9:20	**Erik Mueller**, *IBM*	
9:20 – 9:40	**Mark Finlayson**, *MIT*	
9:40 – 10:00	**Neil Cohn**, *TuftsUniversity*	
10:00 – 10:30	Discussion Panel	**Nicholas Asher**, *University of Texas* **Bart Verheij**, *University of Groningen* **Patrick Winston** moderating + session speakers
10:30 – 11:00	Break	
11:00 – 11:20	**Srini Narayanan**, *University of California at Berkeley and ICSI*	
11:20 – 11:40	**Reid Swanson**, *University of Southern California ICT*	
11:40 – 12:00	**Ken Forbus**, *Northwestern University*	
12:00 – 12:30	Discussion Panel	**Henry Lieberman**, *MIT* **Ian Horswill,** *Northwestern University* **Whitman Richards** moderating + session speakers
12:30 – 2:00	Lunch	
2:00 – 2:20	**Michael Young**, *North Carolina State University*	
2:20 – 2:40	**Pablo Gervás**, *Universidad Complutense de Madrid*	
2:40 – 3:00	Discussion Panel	**Ian Horswill**, *Northwestern University* **Emmett Tomai**, *University of Texas Pan American* **Whitman Richards** moderating + session speakers
3:00 – 3:30	Break	
3:30 – 3:50	**VS Subrahmanian**, *University of Maryland*	
3:50 – 4:10	**Bart Verheij**, *University of Groningen*	
4:10 – 4:30	**Jerry Hobbs**, *University of Southern California ISI*	
4:30 – 5:30	Discussion Panel	**Michael Cox**, *DARPA* **Patrick Winston** moderating + session speakers
5:30 – 7:00	Cocktails	
7:00	Dinner	**Jay Keyser**, *MIT*

Saturday, October 10

8:30 – 9:00	Opening Remarks
9:00 – 10:30	Discussion
10:30 – 11:00	Break
11:00 – 12:00	Discussion
12:00 – 12:30	Closing Remarks
12:30 – 2:00	Lunch

Appendix 2: Participant Abstracts

Nicholas Asher, Department of Philosophy, University of Texas

What is narrative? One possibility is that a narrative is just a text with discourse structure; but texts without events certainly have discourse structure, but clearly stories without events are rare.

I claim narrative is a particular genre or type of discourse. To understand the species, we must understand the genus with the hope of isolating, if not necessary and sufficient properties, at least prototypical properties that individuate narrative from other genres. Discourse is a structured linguistic object. Discourse structure resembles syntactic structure in that it is recursively constructed from what my colleagues and I call elementary discourse units or EDUs, which are linked together by discourse relations. Linked EDUs can serve to construct complex constituents that are also linked by discourse structure. (1) provides an example of the hierarchical structure induced by discourse relations and the recursive construction process.

(1) (a) (π_1) John had a great evening last night.
(b) (π_2) He had a great meal.
(c) (π_3) He ate salmon.
(d) (π_4) He devoured lots of cheese.
(e) (π_5) He then won a dancing competition

Statements (1c-1d) provide 'more detail' about the event in (1b), which itself elaborates on (1a). Statement (1e) continues the elaboration of John's evening that (1b) started, forming a narrative with it (temporal progression). The ordering of events does not follow the order of sentences, but rather obeys the constraints imposed by discourse structure, as shown graphically below. Thus the eventualities that are understood as elaborating on others are temporally subordinate to them, and those events that represent narrative continuity are understood as following each other. A theory like Segmented Discourse Representation Theory (SDRT: Asher 1993, Asher and Lascarides 2003) provides the discourse structure for (1) above in for Figure 9 below, which provides a proper temporal structure for the text and has several other effects on content. The elements π_6 and π_7 in the figure are complex discourse constituents created by the process of inferring the discourse structure (Asher and Lascarides 2003).

Within this framework of a view about discourse, we can ask what is a narrative. It

is also be a particular kind of structured object. But it is not simply a sequence of EDUs linked by Narration, but more something like the structure given in (1). I advocate an empirical approach to an investigation of this genre: by annotating paradigm examples of narratives, I hope that certain general structural patterns and or intentional structures will emerge that will serve to characterize the genre.

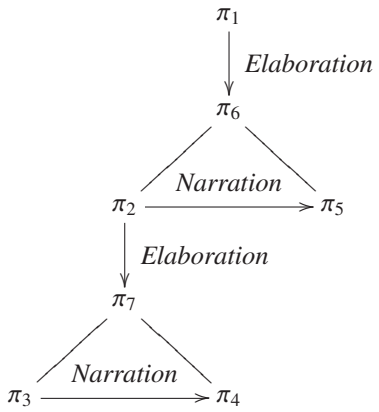

Figure 9: The SDRT structure for (1).

Neil Cohn, Department of Psychology, Tufts University

Foundations for a Visual Language Grammar Drawing joins speech and gestures as one of only three ways in which humans can express concepts. Like those forms, the visual-graphic form can also be put into discrete sequences that requires a degree of fluency for comprehension and production, and emerges in a variety of socio-cultural contexts, from modern day "comics" to native Australian sand narratives. I argue that the understanding of sequential images is guided by a hierarchic generative grammar that is structured with similar principles as language, with grammatical categories, recursion, distance dependencies, structural ambiguities, and grammatical constructions. Throughout, the foundations will be laid for a theoretical model of this "visual language" narrative grammar, establishing the basic issues involved and concerns posed to future research.

Mark A. Finlayson, CSAIL, MIT

Deriving Narrative Morphologies via Analogical Story Merging It has long been suspected that stories drawn from the same cultural setting share important narrative structure. One example of such structure, first identified by Vladimir Propp in 1928, is the morphology of a set of stories, which describes the set of plot elements and their allowed sequences. Until now, the extraction of morphologies has remained a manual task, the purview of anthropological virtuosos (e.g., Propp, Lévi-Strauss, Dundes, Campbell); reproduction or validation of their analyses is a

time-consuming, prohibitively difficult endeavor. I use a technique called Analogical Story Merging that derives a morphology given a set of stories. It incorporates standard techniques for computing analogies as well as a method called Bayesian Model Merging for inducing a grammar from a corpus of positive examples. The output of the basic implementation has been applied to a small example story corpus, a set of summaries of Shakespearean plays.

Kenneth D. Forbus, Department of Computer Science, Northwestern University

Analogy and Narrative: Modeling Human Decision-making A growing body of evidence in cognitive science suggests that analogical processing is one of the core operations of human cognition. Analogy is powerful because it enables particular experiences to be used to reason about, and make decisions in, novel situations, as well as learn generalizations at human-like rates. Narratives can be viewed as a distillation of experience, providing a source of guidance even for situations one has not directly experienced. The key idea involves "structure-mapping" and how it can be used to model aspects of moral decision-making presented as a story. Our work so far suggests that the combination of analogy and narrative may provide a useful new methodology for modeling cultural influences on reasoning. (Joint research with Morteza Dehghani and Emmett Tomai.)

Pablo Gervás, Department of Software Engineering and Artificial Intelligence, Universidad Complutense de Madrid

Story Generation: Composing and Inventing Stories A story can be thought of as a linear discourse that tells about a set of events that take place at more than one location over a number of time periods and involve several characters. If this set were mapped on a space-time grid, it will most probably not look linear at all. When causal relationships between different events are considered, a complex interconnected graph results. A fundamental subtask of the storytelling process is how the non-linear set of events gets mapped onto a linear discourse. Empirical observations of human-made stories indicate that causal relationships play an important role in structuring and organizing the story, but they are often not mentioned explicitly. Rather, they are left for the reader to infer. A very important ingredient when attributing merit to a story seems to be related to the number of such implicit causal relations that can be inferred when reading it. Our research over recent years focuses on three key issues traditionally overlooked by story generators: how stories come to be linear when they refer to complex clouds of events unevenly distributed over space and time (discourse planning), what role does the process of interpretation expected of the reader play in story composition

(reader models), and the development of algorithms capable of generating different new stories in each run (creative storytelling).

Jerry Hobbs, ISI, University of Southern California

Narrative and Planning In the Strong AI paradigm, people are viewed, in part, as planning mechanisms. We have the general goal to thrive, and we continually construct, execute, monitor and modify a plan to achieve that goal, using our beliefs about what kinds of actions tend to cause what kinds of results, including the result of thriving. A narrative is an account of such a planning mechanism attempting to achieve its goals in the face of obstacles. The unique power of narrative derives from the fact that it presents a course of events in very much the way we, as dynamic planning mechanisms, experience them in our own lives. Several very short stories illustrate this perspective. They reveal the central importance of the intersection of ontology and natural language processing, namely, concepts related to causality and how they are realized in language, and concepts from an ontology of micro-sociology, expressing many of our most central concerns.

Ian Horswill, Departments of EECS and Radio, Television, & Film, Northwestern University

Modeling Mammalian Neuropsychology for Interactive Virtual Characters One of the primary differences between humans and other animals is our capacity for high level cognition and the use of language. Historically, AI has understandably focused on duplicating these aspects of human behavior, and selected its architectures accordingly. However, humans are nevertheless social mammals and share a large part of the mammalian behavior repertoire, such as fight, flight, and feeding, as well as social behaviors, such as attachment, affiliation, territoriality, and the formation of dominance hierarchies. I argue that for all our unique capabilities, the mammalian behavior hardware (whatever it may be) is still active in humans, and that for applications such as virtual characters for interactive drama, we should begin from architectures based on our commonalities with other social mammals, rather than from our unique capabilities.

Leroy A. “Jack” Jackson, Naval Postgraduate School

The narrative paradigm serves to harmonize various human cultural, social, and behavioral theories in a cultural geography model under development by the US Army TRADOC Analysis Center and various partners. Walter Fisher's narrative paradigm contends that people are essentially storytellers, that all meaningful communication is a form of storytelling, and that the world is a set of stories from

which each individual chooses those that match his values and beliefs. The Cultural Geography model (Jackson 2009) is a prototype implementation of an agent based model of civilian populations in stability operations derived from social theory and military doctrine. U.S. Army and Marine Corps doctrine claims that the most important cultural form for commanders to understand in a conflict environment is the narrative. A cultural narrative is a story relating a causally linked set of events that explains some aspect a group's history and expresses the group's values, character, or self-identity. Members of a society express and absorb ideologies through narratives. Therefore, by attending to narratives, commanders can identify a society's or group's core values and begin to understand their methods of reasoning and their behaviors.

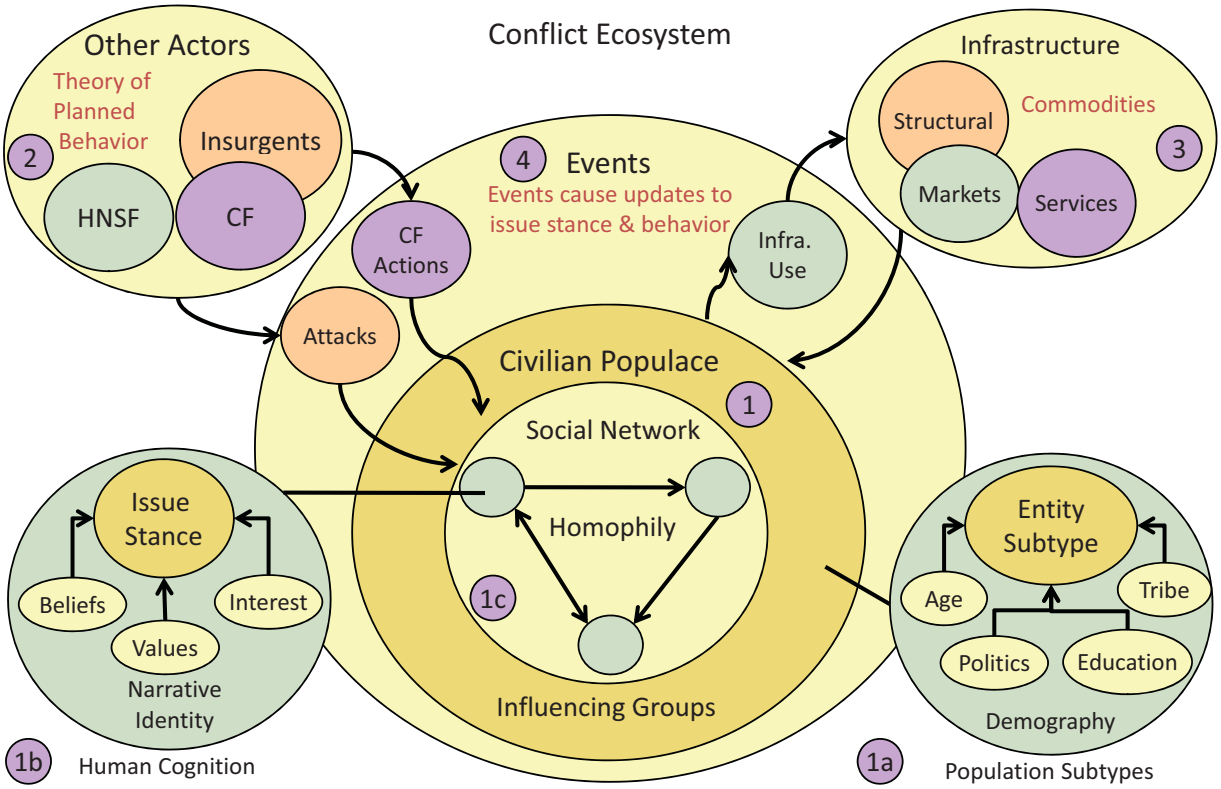

Figure 10: This is a conceptual overview of the cultural geography model and it components. We view the environment as a conflict ecosystem in which both the coalition forces and threat forces are components. The four main components of the model are (1) the civilian populace, (2) other actors within the model, (3) infrastructure objects, and (4) events that occur within the model (4). Narrative informs population entity identity and behavior, and the influence of events on population entities.

Samuel Jay Keyser, Department of Linguistics, MIT

How to Write a Story, Sometimes Edgar Allan Poe's essay "The Philosophy of Composition" begins with a quotation from a note by Charles Dickens:

> *By the way, are you aware that Godwin wrote his Caleb Williams backwards? He first involved his hero in a web of difficulties, forming the second volume, and then, for the first, cast about him for some mode of accounting for what had been done.*

I cannot think this is the precise mode of procedure on the part of Godwin…but the author of Caleb Williams was too good an artist not to perceive the advantage derivable from at least a somewhat similar process. Nothing is more clear than that every plot, worth the name, must be elaborated to its dénouement before anything be attempted with the pen. It is only with the dénouement constantly in view that we can give a plot its indispensable air of consequence, or causation, by making the incidents, and especially the tone at all points, tend to the development of the intention. I will illustrate the truth of these observations with selections from a series of children's poems published under the title The Pond God and Other Stories by Samuel Jay Keyser.

Henry Lieberman, Media Laboratory, MIT

Common sense reasoning for understanding and generating narrative Since the earliest days of natural language understanding, researchers have realized the need for background Commonsense knowledge and Commonsense reasoning to fully understand or generate meaningful narrative content. There's simply not enough information explicit in narrative text to fully convey meaning. For the past ten years, we have been collecting Commonsense knowledge from volunteers over the Web. The Open Mind Common Sense corpus contains over 1,000,000 sentences in English (along with collections in other languages). It is "the Wikipedia version of Cyc". We have applied state-of-the-art parsing and created a semantic net, ConceptNet, that we believe represents roughly 1% of what the average person knows. We have a new reasoning technique, AnalogySpace, that is more suited to Commonsense inference than traditional logical reasoning. It can reason over an entire knowledge space at once, discovering the dimensions that best characterize the space. It can classify concepts according to those dimensions, create categories, or compute the plausibility of assertions. The tools are available open-source. It has been applied to a myriad of tasks involving narrative, such as analysis of user opinions of software, and generative tasks, such as an intelligent video editor that selects clips based on their appropriateness for a story.

Erik T. Mueller, IBM Research

Story Understanding through Model Finding My main interest is machine story understanding and how to build story understanding systems. Although my basic approach – building models of an input story – has remained stable, my method for building these models has changed. Initially, I used specific algorithms for understanding each domain (like space, time, goals, and emotions). Over time, however, I was led to my current approach of using efficient, general algorithms that operate on declarative knowledge. Changing my approach allowed me to scale up my systems from understanding 3 stories to understanding 107 stories and beyond. The key component is the Discrete Event Calculus Reasoner, which builds story models given a commonsense knowledge base and predicate-argument representations of story events and states.

Srini Narayanan, Cognitive Science Program, UC Berkeley and ICSI

Simulation semantics: A computational framework for exploring the links between language, cognition and action. The UCB/ICSI NTL project is an ongoing attempt to model language behavior in a way that is both neurally plausible and computationally practical. Work within the NTL project coupled with a variety of converging evidence from Cognitive Linguistics, Psychology and Neuroscience suggests that language understanding involves embodied enactment which we call "simulation semantics." Simulation semantics hypothesizes the mind as "simulating" the external world while functioning in it. The "simulation" takes noisy linguistic input together with general knowledge and makes new inferences to figure out what the input means and to guide response. Monitoring the state of the external world, drawing inferences, and acting jointly constitute a dynamic ongoing interactive process. My advance is a computational realization of a simulation semantics hypothesis, with preliminary results on applying the model to vexing problems in narrative interpretation. Specifically, the key notion is an implemented computational model of sensory-motor imagination combined with the ability for metaphoric projections which potentially explains the cross-cultural disposition to conceptualize abstract actions and events in terms of sensory-motor representations (consider for example the domain of economic events where economies stumble, lurch, speed up, slow down, sprint, crawl, move-ahead, turn around, reorient, gain a foothold, etc.). Results of the model and ongoing imaging and behavioral experiments will be described along with a new computational formalism called Embodied Construction Grammar (ECG) that supports a deep semantic analysis of narratives by integrating theories from construction grammar, frame semantics, and cognitive linguistics.

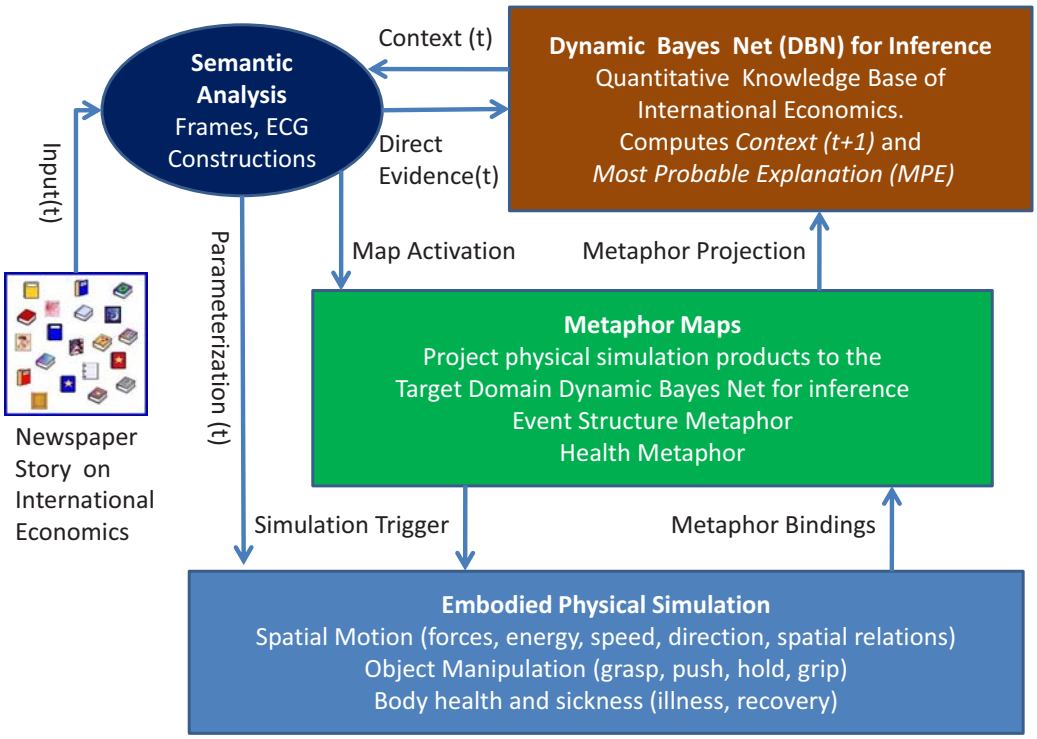

Figure 11: Narayanan's KARMA system exploits metaphoric projections of physical reasoning and simulation to interpret narratives about international politics and economics.

V.S. Subrahmanian, Department of Computer Science, University of Maryland

Generating Succinct Stories Suppose you wanted to query the web with a query which said "Tell me the story of Paul Kagame in 10 lines or less." The answer to this query should depend not only on the topic (Paul Kagame) but also on the interest the user has on the topic. A historian interested in Kagame may be interested in his early childhood, while a State Department African policy expert may have no interest in this. We develop a model of stories that can be extracted from text sources and that are parameterized by user interest. The user specifies his level of interest in various attributes of the topic and the goal of the system is to create a story that (i) maximizes interest to the user, and (ii) minimizes repetition, and (iii) maximizes continuity in the story. The formal framework generates such stories. We also have results of experiments conducted on the utility of the framework. The work was done jointly with M. Albanese, C. Cesarano, M. Fayzullin, and A. Picariello.

Reid Swanson, ICT, University of Southern California

Encoding the common-sense knowledge, in a machine-readable format, that enables the level of reasoning power to support a person's ability to tell compelling stories is a daunting task. In the spirit of Open Mind we believe allowing authors to engineer this knowledge in natural language is a critical component to covering the wide scope of human activities. However, we also believe that much of this knowledge is being spontaneously volunteered to the web in the form of personal weblog diaries. As of 2008 Technorati.com reports that nearly 1 million weblogs are authored every day and over 133 million have been indexed since 2002. From our own research, somewhere between 5-17% of these weblogs are individuals telling personal stories about their lives. These stories are an invaluable resource that contains all sorts of temporal and causal knowledge about peoples' activities, relationships and thoughts. Our system leverages the vast repository of stories found on the web, in a surprisingly simple way, but is still capable of enabling a real-time interactive storytelling system, in which a human and computer collaborate to write a narrative together.

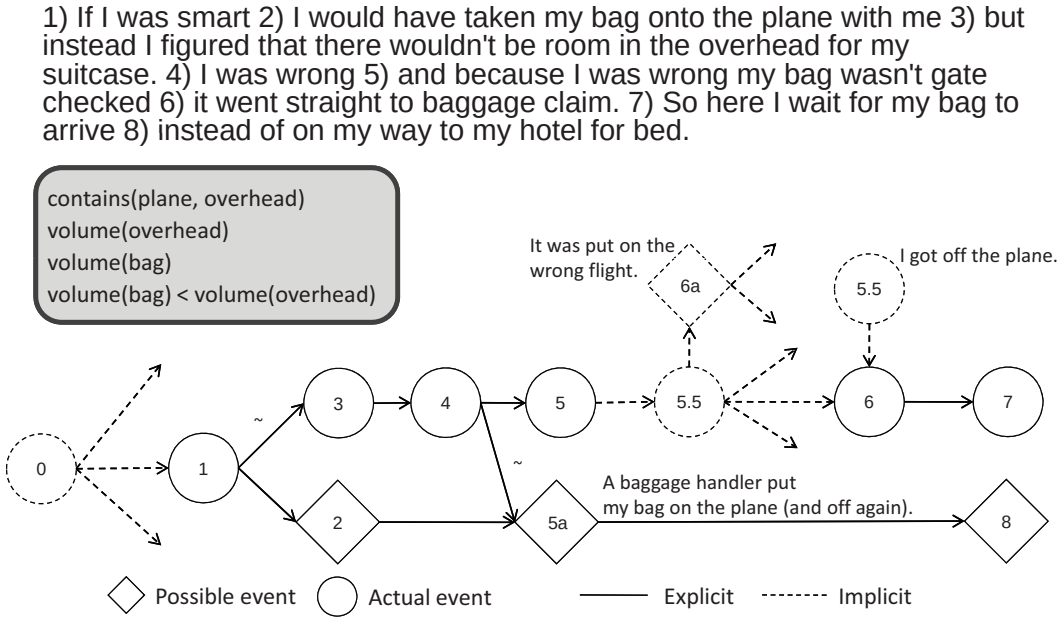

Figure 12: A sample narrative analysis (Swanson 2008)

Emmett Tomai, Department of Computer Science, University of Texas Pan American

Understanding narrative is a hallmark of human communication, and a difficult, long standing challenge for artificial intelligence. Even the question of what it means to understand a narrative has fragmented and often unclear answers. In recent work we have focused on pragmatic reasoning tasks over narrative text to

provide a clear account of some different facets of narrative understanding. We have used cognitive models of decision-making and blame attribution in response to narrative stimuli, as well as a heuristic model of identifying the intended morals of fable narratives. These reasoning tasks provide a well-defined context for exploring narrative representation issues and the impact of pragmatic concerns on the semantic interpretation process. Key challenges include the need for broad commonsense knowledge, highly structured event representations and the necessity of representing nested hypothetical situations (e.g. utterances, possible or conditional futures, etc). We have built a natural language understanding system, EA NLU, that is capable of **g**enerating such knowledge-rich, highly expressive logical representations by using pragmatic constraints to guide and control the interpretation process. EA NLU integrates the ResearchCyc knowledge base contents, compositional frame semantics with delayed disambiguation, discourse representation theory and abductive reasoning. This is joint work with Ken Forbus.

Bart Verheij, Department of Artificial Intelligence, University of Groningen

Argumentation Schemes, Stories & Legal Evidence: A Computational Perspective Imagine yourself in court, having to defend your innocence of a serious crime. Let's suppose that your defense fails, and you end up behind bars. Was it your – probably imperfect – control of formal argumentation techniques that made you lose? Or, was the problem more a matter of the content of the stories you told, such as your unconvincing alibi? The starting point of the project "Making sense of evidence: Software support for crime investigations" was that both argumentative and narrative elements are relevant when deciding about the facts in a criminal case. Whereas arguments allow for the careful assessment of individual pieces of evidence, stories give insight into the quality of the overall picture of a case and help to avoid tunnel vision. A result of the project is a hybrid theory of argumentation with stories, thereby showing that previous argument-based and story-based approaches can be naturally combined. The project has taken inspiration from developments in the design of argumentation support software and from argumentation schemes research in the field of argumentation theory.

Patrick H. Winston, CSAIL, MIT

I believe storytelling and understanding is the defining competence of human-level intelligence, and that competence, in turn, rests on human language. Language gives us words, which enable indexing, and language gives us descriptions, which enable analogy.

There is more to the story, however. In the Genesis project, we aim to understand how descriptions enable what Minsky calls self reflective thinking in his book, *The Emotion Machine*. We also aim to understand how language marshals the resources of perceptual systems that solve problems not readily solved symbolically. And we also aim to understand how language enables us humans to imagine situations we have never witnessed before, so that we can use linguistically and visually imagined situations as surrogates for direct experience. We believe all these abilities are essential not just to storytelling and understanding but also to a full understanding of human intelligence.

In the story understanding dimension, we have built the Genesis system, which features more than a dozen representational experts coupled with the ability to use background commonsense knowledge to fill in gaps. In Figure 13, the Genesis system displays its interpretation of a very brief rendering of the Macbeth story, showing in white boxes that which is explicit in the story, in gray boxes that which was inferred using commonsense knowledge, and in pink and yellow, an instance of a higher-level notion, "revenge," discovered by a program based on Wendy Lehnert's plot-unit research.

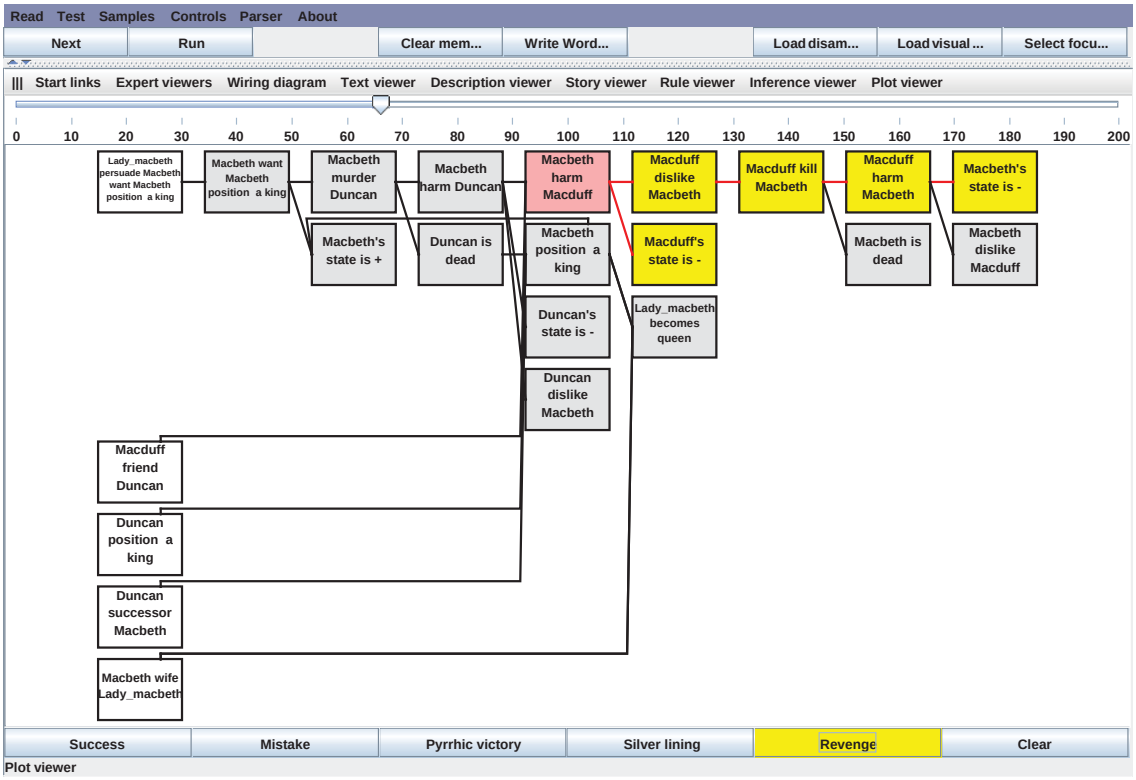

Figure 13: Screenshot of the Genesis System

R. Michael Young, Department of Computer Science, North Carolina State University

Research in the Liquid Narrative Group I seek to build an generative computational model of narrative that is informed both the strong history of prior analytical work in both cognitive psychology and narrative theory. Breaking narrative into distinct story and discourse elements, I build on existing AI approaches to reasoning about action and communication to generate interactions within 3D virtual worlds that should be readily understood as narratives by users. My group has specifically focused on elements of story generation that target narrative phenomena such as character intentionality, suspense, surprise and expectation. At the discourse level, we have developed methods for the automatic generation of cinematic camera control for narratives told in 3D worlds. Our work on interactivity in narrative environments has produced methods for analyzing and dynamically adapting stories in response to user activity. Central to this work are three key concepts. First, that the development of precise, generative computational models of narrative must be informed by ideas drawn from a range of inherently non-computational, semi-formal disciplines. Second, that narrative is ultimately in the head of the consumer – that is, the ultimate design criteria for narrative artifacts rest in the cognitive and affective responses they prompt in their human consumers. And finally, existing models of narrative must be extended and adapted to account for new modes of narrative in which user interactivity takes center stage.