Narrative Presentation and Meaning

Emmett Tomai and Kenneth D. Forbus

Qualitative Reasoning Group, Northwestern University
2133 Sheridan Road, Evanston, IL, 60208, USA
{etomai, forbus}@northwestern.edu

Abstract

We present a preliminary sketch of our exploration of understanding event descriptions by casting them in terms of narrative presentation. We maintain that human use of narrative in communication reflects a consistent approach to the organization of information. We further hypothesize that this method of organization is an effective heuristic for focusing inference to understand the world around us. Thus, we believe that casting events and descriptions of events in terms of narrative presentation will both increase the salience of mechanical understanding results and decrease the computational complexity.

Introduction

In a recent National Basketball Association game, a defensive player, we'll call him "B", kicked an offensive player, we'll call him "A", in the back of the leg as he jumped up to dunk the basketball. After the game, player A publicly stated that the kick was an intentional attempt to injure him and that player B has a history of trying to injure people. Videos of the incident and numerous others involving the same player were quickly uploaded to YouTube, slowed down, zoomed in and analyzed on sports message boards around the internet. From that same set of events – a simple physical sequence in a well understood context – It was argued by numerous fans that the kick was intended to injure, that it was a common defensive "trick" intended to throw the other player off, that it was purely accidental, that player B actively tries to injure people, that player A made a mistake by talking about it, that his team needs to retaliate and so on and so forth.

Narrative is interesting. That is to say, the point of narrative presentation is that it serves to extract and highlight those events, outcomes and conclusions which the presenter and observer find "interesting". While this is a statement about the art of story telling, it also highlights one of the most fascinating properties of human reasoning over everyday events – the ability to observe a situation and quickly identify a small set of salient implications out of a massive space of possibilities. We believe that the use of narrative as a communication medium reflects a structured approach to reasoning over this space and that a computational model that aspires to understand in a human-like way would be well informed to follow that approach. In this paper we present a framework for approaching the problem of understanding events and the stories that describe them. We discuss first what makes a story and what it means to understand it.. We then describe the approach that we are exploring to create such a process. Finally we describe experimental applications that we are working towards and future directions.

Understanding stories

While there are many definitions of story to choose from, we opt for a simple working definition: the presentation of a sequence of observed events. The events that the story is based on may be real or imagined, physically observable or entirely mental. This narrative presentation will always involve a subset of partial representations of the source events combined with additional assertions. Although we are not particularly concerned here with delineating what is and isn't a story, we do assume that the presenter has some set of communication goals and that the elements of the presentation are coherent to those goals. That is, that each utterance in the presentation adheres to Grice's maxim of relevance (Grice 1975). If the subset is the entire set, the representations are complete and no additional assertions are made, then we have a literal presentation of the events, such as a still video camera might capture. This has no goals and thus is not a story we are interested in, but rather a boundary condition, one end of a spectrum which shows more and more *interpretation* as it moves away.

A vast amount of research in different fields has been done which bears on the idea of story understanding. We present here several particular threads of research that have contributed to our goals and approach.

Artificial intelligence researchers have directly attacked in-depth semantic understanding of narrative texts for decades. In the 1970s Charniak (Charniak 1972) tackled the problem of understanding children's stories while Schank (Schank 1977) constructed a theory of knowledge systems based on understanding textual stories. This theory produced several notable story understanding systems including BORIS (Dyer 1983) which integrated much of the previous work covering scripts, goals and plans with themes. McCarthy's memorandum on natural language understanding (McCarthy 1990), first written in

1976, defined the understanding task in terms of the ability to answer a wide range of deep questions about what had taken place in the story. This competence involves identifying unstated events and effects, identifying events and effects which had some possibility of happening but did not, explaining causal chains, predicting future events and effects and evaluating events and effects in terms of utility as well as values. A 1999 update by McCarthy observed that the problems he presented had not yet been solved. Mueller (Mueller 1999) has taken up that torch by identifying and attacking each necessary area of competence with domain specific agents.

Much of this research has served to demonstrate the many areas of reasoning competence required for the task. Stepping back from the difficulties of natural language processing, there are layers of complexity in reasoning over the most commonplace events and states. Notable here are advances in hierarchical planning and plan recognition. Further, this is complicated by introducing intentional agents that require inferences about unobserved beliefs, goals and intentions (c.f. Bratman 1987). Leaving the realm of purely physical interaction and moving to social interaction introduces more general character attribution where internal states such as emotions, value systems and personality tendencies are inferred from external actions (c.f. Kelley 1973, Gratch, Marsella & Mao 2006). It also makes unavoidably clear the fact that real-world events are multi-purpose, playing different roles based not only on context but point of view.

Reasoning about the implications of anything but the simplest, most constrained physical events presents a practically infinite space of possible conclusions and predictions. The most straightforward model of understanding - one which generates a single interpretation that makes all the relevant inferences and can thus explain, predict and evaluate completely - is neither tractable nor sensible. As previously stated, we believe that narrative serves as both a communicative and organizational device to encode a particular subset of those inferences – a particular *meaning* among many possible interpretations. An obvious implication of this position is that one attempting to understand the content of a narrative while ignoring the presentation will result in unnecessary complexity explosion. To put it positively, narrative presentation serves to constrain complexity and guide inferential attention.

Turning then to work in natural language understanding, evaluation is a different story. Few projects approach in-depth semantic understanding but rather focus on building blocks. Disambiguation of word senses, identification of entities and events, anaphora resolution, role and preposition attachment and textual entailment are broadly covered. Deeper semantic concerns can be likewise approached first in terms of temporal and causal networks. Within this reasonably tempered approach there is solid progress on the application of presentation elements for constraint. Discourse Representation Theory (DRT) (Kamp and Reyle, 1993) demonstrated how ambiguities over quantification and tense/aspect could be resolved by paying attention to clausal structure. Segmented DRT (SDRT) (Asher and Lascarides 2003) takes this a step further by showing that proper identification of rhetorical relations between clauses (Hobbs 1985, Mann and Thompson 1987) is necessary for dealing with certain lexical, syntactic, semantic and pragmatic ambiguities. SDRTs principle of Maximize Discourse Coherence serves as an effective heuristic for constraining ambiguity through discourse-level relevance.

However, narrative is not merely coherent, it is directed.

Narrative presentation

Narrative presentation is ubiquitous in human communication. In spite of the subjectivity of interpretation, it is clear that narrative presentation is an effective and reliable, even preferred, way of conveying information. Some are better at it than others, but overall human competence is quite high. How is it, with so many possible interpretations, that a presenter can reliably convey the one he or she intends? Shared context and assumptions are critical (c.f. Stalnaker 1998); we propose that shared narrative devices are likewise.

Beyond coherence and relevance, a skilled narrator will rely on narrative devices to cue the interpretation he or she is attempting to convey. As we saw in our example, there is no shortage of different interpretations over observable events. What is interesting is that so many of these divergent interpretations generated significant support as plausible explanations of the events. We hypothesize that for a given corpus of narrative presentations, a set of narrative devices can be identified which can serve to heuristically focus inferential attention resulting in interpretations that match human criteria for plausibility and relevance.

There are numerous theories of narrative structure that have been explored in several fields. We present again a selection that has influenced our direction.

In the field of social psychology, research on narrative comprehension has seen the development of *story grammars* as a mental organization of events and states (c.f. Mandler & Johnson 1977) and *causal network theory* highlighting the importance of causal connections in recall and summarization tasks (Trabasso, et. al 1984, van den Broek & Lorch 1993).

Several approaches in linguistic studies and literary criticism have approached the question of narrative in terms of functional parts. Vladimir Propp's analysis of folk tales (Propp 1969) presented a coarse-grained structure of narrative while William Labov's analysis of oral storytelling (Labov 1966) suggests a function of *evaluation* which indicate to the listener which parts of the story are relatively more important, resulting in inferences about the storyteller's attitude towards the story. Polanyi drew from Labov as well as several other sources to apply a broad battery of evaluative devices at the syntactic, semantic and discourse levels (Polanyi 1989). Roland

Barthes (Barthes 1977) makes two observations on the nature of functional narrative units. The first is a distinction between *cardinal* units that are hinge-points of the narrative plot and *catalysers* that cluster around them, filling the space between them with reasonable, relevant activity. The second observation is that narrative units may simultaneously belong to multiple classes. Beyond cardinal and catalyzing functions, Barthes speaks of the important class of *indices* that indirectly refer to diffuse concepts such as character and atmosphere; ideas that contribute to the meaning of the narrative even though they have no bearing on the events taking place.

Returning to natural language studies in artificial intelligence, Asher and Lascarides demonstrate the effective application of 36 rhetorical discourse relations including *narrative*, *contrast*, *background* and *elaboration* (Asher & Lascarides 2003).

In our exploration we have been building on these foundations to consider such devices as conflict and resolution, fulfilling or violating expectations and assumptions, establishing questions and revealing answers, and implying causality. Another class of devices pertains to the issue of human character: roles and archetypes can be applied and recognized while opportunities for specific personality attributions are provided by reactions to challenges, forced choices and the process of character development. Lastly, we are investigating devices for directing value judgments: focus on character choices as well as more abstract value-laden concepts such as justice, sacrifice, honor and dedication.

Experimental directions

There are three areas in which we intend to utilize and evaluate an implementation of our model of understanding. The first experimental task is extracting temporal and causal networks from a corpus of Aesop's Fables, demonstrating the effect of narrative presentation structure as a constraint on inferential attention. The second task is to take different interpretations of the same (real or imagined) source events and demonstrate that subtle differences in presentation are captured by our narrative structure and result in different inferences. We are particularly interested in the expectations raised by each interpretation as predicting salient future outcomes. The third experiments are part of a joint work with Doug Medin's psychology group at Northwestern University to explore the cognitive processes underlying moral reasoning. Medin's group is investigating how people from different cultures reason about moral dilemmas, particularly with respect to *sacred values* vs. traditional consequentialist reasoning (Tanner et. al, forthcoming). Our goal is to create computational models of such decision-making. We hypothesize that our model of understanding will be able to identify from the presentation of the dilemmas which elements of the situation respondents will find salient in the decision. Those elements will then be able to inform our decision-making model to identify protected values and predict responses. We will also use our sequential generalization engine SEQL (Kuehne, *et al* 2000) based on Gentner's structure-mapping theory (Gentner 1983) to predict the responses of others in the same cultural group to test the hypothesis that the interpretation of meaning is culturally informed.

In both experimental setups, we will take as input natural language text using a semi-automatic tool for efficient knowledge acquisition. The source material is manually translated into a controlled language which is then semi-automatically translated into formal representation. We are using the Explanation Agent NLU system (EA), originally constructed by Kuehne (Kuehne 2004) as part of an exploration of how qualitative process theory (Forbus 1984) could be used in natural language semantics. Kuehne demonstrated that EA could automatically generate QP descriptions from controlled language text. In order to broaden EA to the ambiguity of open ended stories we are working in collaboration with Mark Finlayson and Patrick Winston at MIT to build a workbench to support human interaction with the system in semi-automatic mode. The workbench provides real-time feedback for the controlled language encoding and tools for disambiguating the formal representations. This is inspired by both CMU's KANT project (cf. Nyberg *et al* 2002) and Boeing's controlled language work (cf. Clark *et al* 2003). For more information on these ongoing projects, see (Forbus et al, 2007).

Related Work

Recent work in understanding narrative has been done by Moorman (Moorman 1997) with the ISAAC system. ISAAC identifies a set of supertasks for more robust story understanding, one of which is story structure including such concepts as character roles, setting, genre and plot summarization. Mueller's ongoing work (Mueller 2004) uses a satisfiability solver to construct multi-representation models across realms such as space, time, needs and feelings.

The recent surge of interest in interactive forms of narrative has seen focus on the role of narrative structure in generating stories. The IDTension system (Szilas 2003) guides the ongoing narrative based on evaluating how satisfying the succession of actions is. The function for satisfaction includes such narrative devices as relevance, ethical consistency and conflict. The Façade interactive drama (Mateas 2003) is built around the concept of dramatic beats which change dramatic values such as love, trust and tension. Search-based drama management (SBDM) (Weyhrauch 1997) uses an author-specific evaluation function to capture the aesthetic that the author intends the story to carry.

Acknowledgements

This research is supported by a MURI grant from the Air Force Office of Scientific Research.

References

Asher and Lascarides. 2003. *Logics of Conversation*. Cambridge University Press, Cambridge, UK.

Barthes, R. 1977. Introduction to the structural analysis of narratives. *Image, music, text*; essays selected and translated by Stephen Heath. Hill and Wang, New York.

Bratman, M. 1987. *Intention, Plans, and Practical Reason*. Harvard University Press. Cambridge, MA, USA. 1987.

Charniak, Eugene. 1972. *Toward a model of children's story comprehension* (AI Laboratory Technical Report 266). Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Clark, P., Harrison, P. and Thompson, J. 2003. A Knowledge-Driven Approach to Text Meaning Processing. *Proceedings of the HLT Workshop on Text Meaning Processing*, pp 1-6.

Dyer, M. 1983. *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*. MIT Press Cambridge, MA, USA.

Forbus, K. 1984. Qualitative process theory. Artificial Intelligence, 24, 1984.

Forbus, K., Lockwood, K., Tomai, E., Dehghani, M. and Czyz, J. 2007. Machine Reading as a Cognitive Science Research Instrument. *AAAI Spring Symposium on Machine Reading*. Stanford University, California.

Gratch, J., Marsella, S., and Mao W. 2006. Towards a Validated Model of "Emotional Intelligence." *Twenty-First National Conference on Artificial Intelligence (AAAI06)*.

Gentner, D. 1983. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.

Grice, H.P. 1975. Logic and conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics III: Speech Acts*, pages 41-58. Academic Press, New York, NY.

Hobbs, J. R. 1985. On the Coherence and Structure of Discourse. Technical Report CSLI-85-37, Center for the Study of Language and Information, Stanford University.

Kamp, H. and Reyle, U. 1993. *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language*. Kluwer Academic Dordrecht; Boston.

Kelley, H. H. 1973. The Processes of Causal Attribution. *American Psychologist*, 28:107-128.

Kuehne, S., Forbus, K., Gentner. D., & Quinn, B. 2000. SEQL: Category learning as progressive abstraction using structure mapping. *Proceedings of the 22nd Annual Meeting of the Cognitive Science Society*.

Kuehne, S. E. 2004. On the Representation of Physical Quantities in Natural Language Text. *Proceedings of the Twenty-sixth Annual Meeting of the Cognitive Science Society*, Chicago, Illinois, USA, August.

Labov, W. and Waletzky, J. 1966. Narrative Analysis: Oral Versions of Personal Experience. *Essays on the Verbal and Visual Arts*, June Helms (ed.) 12-44. Seattle: University of Washington Press.

Mandler, J.M. and Johnson, N.S. 1977. Remembrance of things parsed: Story structure and recall. *Cognitive Psychology*, 9, 111-151.

Mann, W.C. and Thompson, S.A. 1987. Rhetorical Structure Theory: A Framework for the Analysis of Texts. *International Pragmatics Association Papers in Pragmatics*, 1, 79-105.

Mateas, M and Stern, A. 2003. Integrating plot, character and natural language processing in the interactive drama Façade. *1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment.*

McCarthy, John. 1990. An example for natural language understanding and the AI problems it raises. In John McCarthy, *Formalizing common sense*. pp. 70-76. Norwood, NJ: Ablex.

Moorman, Kenneth 1997. *A functional theory of creative reading: Process, knowledge, and evaluation* (Doctoral dissertation). Atlanta, GA: College of Computing, Georgia Institute of Technology.

Mueller, E. 1999. Prospects for in-depth story understanding by computer. arXiv:cs.AI/0003003

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

Nyberg, E., Mitamura, T., Baker, K., Svoboda, D., Peterson, B., and Williams, J. 2002. Deriving Semantic Knowledge from Descriptive Texts using an MT System. *Proceedings of AMTA 2002.*

Polanyi, L. 1989. *The American Story*. MIT Press. Cambridge, MA, USA.

Propp, V. 1969. *Morphology of the Folktale*. Trans. Laurence Scott. Ed. Louis A. Wagner. 2nd edition. Univ. of Texas Press.

Sacerdoti, E.D. 1977. *A Structure for Plans and Behavior*. American Elsevier, New York.

Schank, Roger C., & Abelson, Robert P. (1977). *Scripts, plans, goals, and understanding*. Hillsdale, NJ: Lawrence Erlbaum.

Stalnaker, Robert. 1998. On the Representation of Context. Journal of Logic, Language, and Information 7:3-19.

Szilas, N. 2003. IDtension: a narrative engine for Interactive Drama. *1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment.*

Tanner, C., Medin D. L., Iliev R. (Forthcoming) "Influence of Deontological vs. Consequentialist Orientations on Act Choices and Framing Effects: When Principles are more Important than Consequences".

Trabasso, T., Secco, T. and van den Broek, P.W. 1984. Causal cohesion and story coherence. In H. Mandl, N.L. Stein and T. Trabasso (eds.). *Learning and comprehension of text*. 83-111. Lawrence Erlbaum Associates, Hillsdale, NJ.

van den Broek, P.W. and Lorch, R.F. Jr. 1993. Network representations of causal relations in memory for narrative texts: Evidence from primed recognition. *Discourse Processes*, 16, 75-98.

Weyhrauch, P. 1997. Guiding Interactive Drama. Ph.D. Dissertation, Tech report CMU-CS-97-109, Carnegie Mellon University.