DramaBank: Annotating Agency in Narrative Discourse

David K. Elson

Columbia University, New York City
dke4@columbia.edu

Abstract

We describe the Story Intention Graph, a set of discourse relations designed to represent aspects of narrative. Compared to prior models, ours is a novel synthesis of the notions of goal, plan, intention, outcome, affect and time that is amenable to corpus annotation. We describe a collection project, DramaBank, which includes encodings of texts ranging from small fables to epic poetry and contemporary nonfiction.

Keywords: narrative; discourse; SIG

1. Introduction

We propose in this paper a set of discourse relations designed specifically for modeling narrative discourse. The schemata is called the Story Intention Graph or SIG, and a particular annotation of a narrative is called a SIG encoding. In later sections, we describe DramaBank, a corpus of SIG encodings we have collected with trained annotators.

The SIG is inspired by prior approaches to discourse such as Rhetorical Structure Theory (RST) (Mann and Thompson, 1988) and the Penn Discourse Treebank (Prasad et al., 2008). While some descriptive models such as these have touched on aspects of narrative, including time and modality (Pustejovsky et al., 2003) and sequences of actions (Chambers and Jurafsky, 2008), we attempt here to capture coreference, time, modality, and agency in a unified set of relations geared toward narrative. Agentive relations in particular have been described by cognitive psychologists and literary theorists as key to narrative comprehension; our approach covers several of these, such as those between an action and the intention of its agent (Bundgaard, 2007), between a goal-driven action and its outcome (van den Broek, 1988), between a goal and its subgoal or superordinate goal (Graesser et al., 1994), between an event and an affectually impacted agent (Stein et al., 2000), and between the story-world and the surface discourse as related by a narrating agent (Bal, 1997). Following these examples, we define narrative discourse as having two or more temporally sequential events that are causally connected; we are particularly interested in highly "plotted" narratives with intentional agents, goal-directed behavior, and affectual consequences, among other factors.

In a previous LREC paper (Elson and McKeown, 2010), we described the first phase of this project: 40 encodings annotated according to those relations that describe the text in terms of its temporal structure and predicate-argument structures (propositional modeling). We also described SCHEHERAZADE, a publicly available annotation tool.[1] The current work extends this effort in several significant ways: First, we have introduced the notion of agency to the SIG, and extended the tool to allow annotators to model goals, strategies, attempts and outcomes. Second, we have run a new collection project to elicit such enhanced encodings from annotators, using not only fables but longer and more varied texts. We present the SIG and DramaBank, now consisting of 110 encodings, as a methodology and the beginning of a shared corpus from which we may pursue data-driven investigations of narrative structure.

2. Annotation Methodology

Narrative is an interplay between the minds of agents, the actions they take, the events which befall them, and the perception and transmission of that content in a communicative artifact (Ryan, 1991). The SIG is a schemata for representing narrative that reifies these elements as nodes (entities with coreferent mentions) and arcs (relations) in a semantic network. This exposes coherence on local and global levels: what events happen, when, why, and to whom.

An encoding consists of three interconnected sections called *layers*: the **textual** layer, which represents spans of the original discourse; the **timeline** layer, which represents events and statives that occur in the story being narrated; and the **interpretative** layer, where nodes represent goals, plans, beliefs, affectual impacts, and the underlying intentions of characters (agents) as interpreted by the story's receiver.

In the textual layer, the discourse is divided up into fragments, typically of clause or sentence length. Each fragment is represented by a **Text (TE)** node. Text nodes are chained together by **followed by** arcs so that the order of nodes in the chain reflects the order in which the fragments appear in the original discourse. Text nodes do not need to cover every clause in the discourse, but each Text node

[1] `http://www.cs.columbia.edu/~delson`

A Lion watched a fat Bull feeding in a meadow, and his mouth watered when he thought of the royal feast he would make, but he did not dare to attack him, for he was afraid of his sharp horns.
Hunger, however, presently compelled him to do something: and as the use of force did not promise success, he determined to resort to artifice.
Going up to the Bull in friendly fashion, he said to him, "I cannot help saying how much I admire your magnificent figure. What a fine head! What powerful shoulders and thighs! But, my dear friend, what in the world makes you wear those ugly horns? You must find them as awkward as they are unsightly. Believe me, you would do much better without them."
The Bull was foolish enough to be persuaded by this flattery to have his horns cut off; and, having now lost his only means of defense, fell an easy prey to the Lion.

Table 1: "The Wily Lion", from Jones (1912).

must relate to a node representing deeper meaning in another layer (in particular, a timeline node that itself relates to an agentive interpretation of the narrative).

Events and statives that occur in the story-world, as opposed to fragments of the story's telling, are represented as nodes in the timeline layer. These **Proposition** (P) nodes are temporally arranged by means of *intervals* in the tradition of Allen's work on temporal reasoning (Allen and Ferguson, 1994). Points in a linear timeline are represented as **State** nodes; events connect to these via **begins at** and **ends at** relations (or are left unbounded). States are ordered, $s_1..s_n$. Text nodes connect to equivalent Proposition nodes with **interpreted as**. Because of this dichotomy, we can represent disfluencies in narration when "story time" and "telling time" diverge, such as during flashbacks.

An example SIG encoding for part of "The Wily Lion" (Table 1), a short fable attributed to Aesop, is shown in Figure 1(a). Three Text nodes highlight passages of the source text that are determined to describe story-world events; related Propositions are linked to State nodes which indicate their temporal sequencing. The State nodes connect to a **Timeline** node which represents a modality within the story-world—in this case, "reality," as the events actually occur in the story-world. Other modalities, such mistaken beliefs about the past, can be assigned to states that belong (via an **in** relation) to alternate Timeline nodes.

The interpretative layer represents the receiver's cumulative situation model (Zwaan and Radvansky, 1998) over the course of comprehending the narrative. This layer includes both content that is directly stated in the discourse and content that the annotator infers. Its purpose is to relate timeline-layer and textual-layer content by motivational, intentional and affectual connections, as opposed to temporal connections.

There are four types of interpretative-layer nodes:

Interpretative Proposition (I). An event or stative about some aspect of the story-world featuring a discrete agent. The *content* of an I or P node is a knowledge representation of the aspect in question, such as a predicate-argument structure.

Belief (B). A belief node acts as a frame, inside of which the content (other nodes and frames) is understood to be a state of the story-world in the mind of a discrete agent. For instance, an I node that relates to a Belief node with an *in* arc is a belief in that node's content on the part of the agent. This agent can be a single conscious entity or a set of entities who share a belief.

Goal (G). A goal node acts as a frame for other interpretative content, similar to a Belief. The difference is that the content of a Goal frame is understood to be the state of the story-world as *desired* by the discrete agent or set of agents.

Affect (A). An Affect node represents a positive (beneficial) affectual impact with respect to an agent.

Agency frames—goals and beliefs—can be nested indefinitely to model theory-of-mind interpretations of narrative meaning (Palmer, 2007). For instance, Alice may want Bob to believe that Alice believes that Bob has some property. When an I node is not placed inside an agency frame, it represents content that the narrating agent of the story asserts to be true in the scope of the story-world. (In some cases, though, the story narrator may be unreliable. For stories within stories, inner narrators are themselves agents.)

Propositions in the Timeline layer are connected to interpretative-layer frames and nodes through six arcs: **interpreted as**, **implies**, **actualizes** and **ceases** indicate a functional relationship; **attempt to cause** and **attempt to prevent** indicate agent intention. The difference between the first four is one of directness. *Interpreted as* indicates direct equivalence, *implies* indicates obvious entailment; *actualizes* indicates a positive but indirect causal relationship; *ceases* indicates a negative but indirect causal relationship.

Finally, the affectual impact of an interpretative node can be indicated through the combination of an Affect node (which indicates a particular agent) and a **provides for** or **damages** arc (which indicate positive or negative impact, respectively). For instance, if a P node *actualizes* an I node node which has an outgoing *provides for* arc to an agent's Affect node, this signifies that the events of P positively impact the agent.

Figure 1 shows a possible interpretative encoding for a small section of the "Wily Lion" timeline. The action at s_1, in which the lion watches the bull feed in the meadow, is positively linked to two nodes: the frame indicating that the bull desires to feed, and the feeding action itself (indicating that the bull does, in fact, feed). The feeding action

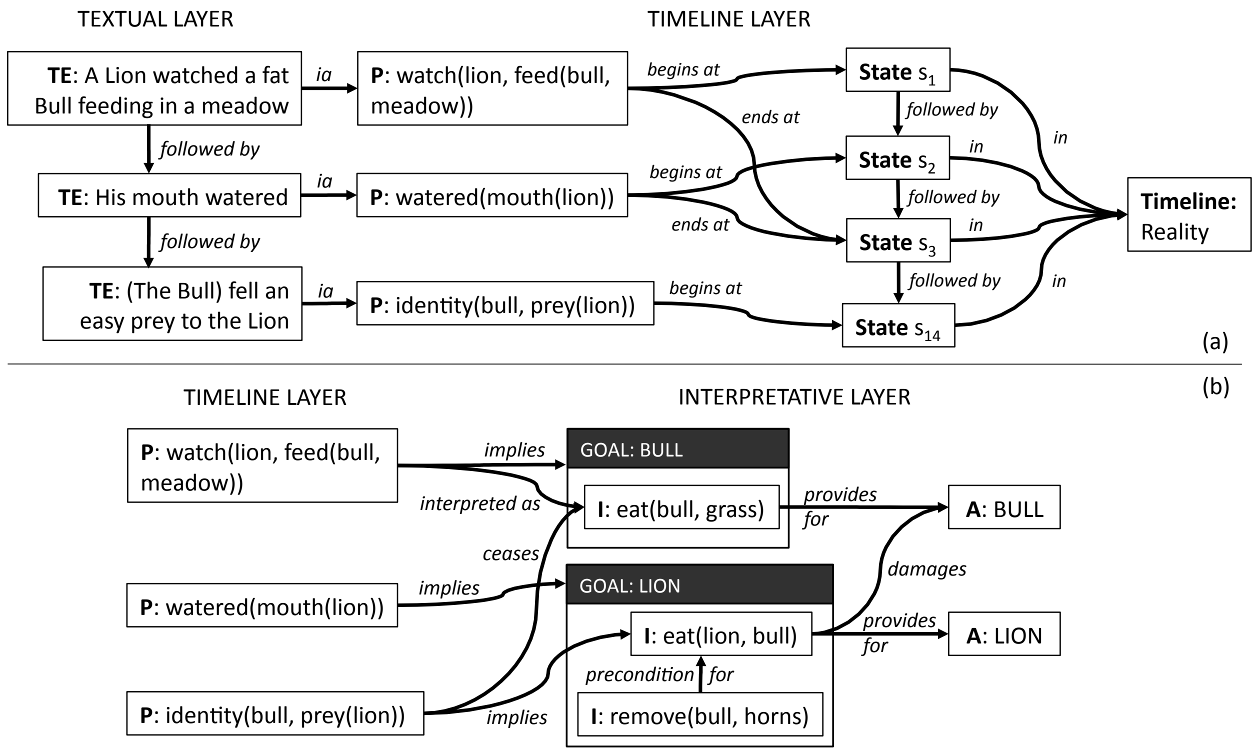

Figure 1: Example SIG encoding of a non-contiguous fragment of "The Wily Lion".

is linked to an Affect node for the bull with *provides for*, a positive association. Thus, the initial story state at s_1 is one in which there is one story-world goal in which the bull desires to undertake a certain action for his own benefit, and the goal is satisfied in that the bull is performing the same action. The "watered" action at s_2 implies that the lion has conceived of a goal to eat the bull; the P node at s_{14} links positively to the lion's goal content (eating the bull, a successful outcome for the lion) and negatively to the bull's goal content (a loss for the bull). The lion has satisfied his goal and provided for his own well-being at the expense of the bull's well-being.

This example also shows how the notion of a linear flow of story-world time is provided to interpretative-layer nodes by means of the timeline layer. For each State node in the "reality" timeline, a set of logical entailments determines which nodes in the interpretative layer are occurring at the corresponding point in story time, and which are not. This computation is called **actualization**. In general, a node's **actualization status** relative to some point in story time is always one of three conditions that describe the objective truth of the node's content at that state:

1. **Hypothetical (H).** The node's content is a hypothetical which may or may not come to pass. The truth of a such a node at the state in question is indeterminate; no assertion is made about whether or not the content is true within the story-world at that moment. All I nodes are hypothetical until changed by an incoming positive or negative arc.

2. **Actualized (A).** The node's content is true (in effect; currently occurring in the story-world). In Figure 1, the bull's goal frame is actualized at s_2, asserting that the bull conceives of the desire at that time; the *eat* node is actualized at s_{14}, indicating success at that time.

3. **Prevented/Ceased (P).** The node's content is false (not in effect; decisively incompatible with the story-world). Nodes that are prevented/ceased not only are untrue at the present time, but given the current state of affairs, have been prevented from happening in the foreseeable future. In Figure 1, the bull's goal content is ceased at state s_{14}.

It is important to emphasize the distinction between actualizing a goal or belief *frame* and the goal or belief *content* inside the frame. Frames themselves refer to the intentional states of their agents; nodes within frames can be actualized or ceased without affecting the actualization statuses of their frames. In Figure 1, the bull's *desire* to eat grass is still actualized at s_{14}, but his goal *content* is ceased, representing a loss. If the frame had instead been ceased at s_{14}, the encoding would instead describe a change of mind.

The "attempt to cause" and "attempt to prevent" arcs indicate the intentionality of agents as they act in the timeline layer. An unintended consequence is one that is actualized by a P node which also connects a separate, intended consequence via one of these two arcs.

Various other narrative scenarios can be expressed through permutations of this vocabulary, such as deliberate harm (one agent acting with intention to harm another, and succeeding), backfire (an agent failing to reach its goal, but unintentionally triggering a side effect that harms it) and comeuppance (an agent harming another intentionally at one state, and later, in an unintended consequence, becoming itself harmed). This permutability is shared between the SIG and Lehnert's influential system of plot units (Lehnert, 1981; Goyal et al., 2010). Following Lehnert, we have cataloged a series of 80 such patterns of relations that represent thematic tropes (Elson, 2012b). As compared to plot units, the SIG relations can express time and intention with greater granularity.

A plan is modeled as a chain of connected nodes inside a Goal frame. Each node acts as a "subgoal" that leads to the ultimate goal at the end of the chain. The connections are directed arcs that indicate expected causality: The agent believes that one subgoal would lead to the superordinate goal that lies next on the chain, and so on, leading to the ultimate goal. Specifically, a **would cause** relation traverses from one interpretative frame or proposition to another interpretative frame or proposition. It signifies that in the belief context of the originating node, an actualization of the originating node would causally lead to (is both necessary and sufficient for) an actualization of the destination node. **Would prevent** is its complement, signifying a belief that the actualization of the originating node would cause the destination node to be prevented/ceased. Two other relations, **precondition for** and **precondition against**, signify a belief that actualization of the originating node is necessary but not sufficient for the actualization or prevention/cessation (respectively) of the destination node. This schematic bears resemblance to a partial-order plan, the key difference being that a SIG plan is an annotator's interpretation of a narrated agent's intentions (Suh and Trabasso, 1993), rather than a solvable system guided by commonsense reasoning.

Figure 2 shows an encoding of the lion's plan. Unlike in Figure 1, the nodes and frames are shaded and marked with their actualization statuses with respect to the moment the lion conceives of his plan—white (/H) for hypothetical, grey (/A) for actualized, and black (/P) for prevented/ceased. At this moment, the bull is dangerous, a fact that is preventing the lion from being able to kill and eat it (a goal which would benefit the lion). The lion's goal is to have the bull remove his horns, which would cease the danger, and, we may entail, restore the lion's ability to kill the bull. The lion's plan is to cause the bull to form a plan of his own, namely, to remove his horns for purposes of vanity. Note that the identical "remove" event appears twice, once in each plan; when the bull does remove its horns, this timeline event actualizes both nodes and furthers both plans (a hidden-agenda pattern). In this manner, the annotator describes the cohesion of the story by linking intention, action, and outcome. The goal node that is linked to the story's beginning (`eat(lion,bull)`) is also linked to lion's plan in the story's midsection, and then actualized as an outcome at its climax. In particular, all of the lion's speech actions are *attempts to cause* the `flatter` action at the head of the plan.

The result is a cohesive model of the fable that integrates its textual, temporal and agentive aspects without relying on world knowledge or a rigidly prescriptive model of discourse structure such as a grammar. In general, the SIG is defined only in terms of its node and arc types, and does not dictate a particular type of knowledge representation for the content of each I and P node (despite the node names). We have used propositions in this example, but the schemata only requires that the annotator give the the identity of the agent relating to each I and P node—the content within these nodes can otherwise adopt an alternate representation scheme, or be left blank.

3. Collection

The DramaBank corpus consists of 110 encodings divided into three collections. The collections differ by their source texts and by the subset of SIG relations that were employed. Collection A, which we described previously (Elson and McKeown, 2010), consists of 40 encodings of 20 short fables attributed to Aesop, but only including the textual and timeline layers. Collection B, using 26 of Aesop's fables, contains 60 complete encodings that feature both propositional modeling (as I and P content) and agentive modeling (interpretative nodes and arcs). For Collection C, annotators created complete encodings without modeling propositions. This was done to accelerate the encoding process and allow for longer and more varied texts to be annotated in terms of their interpretative connections. For the same reason, annotators only encoded as Text nodes those passages in the source texts that related to the interpretative layer; digressions that did not relate to goals, plans, intentional actions, affectual impacts or outcomes were left out of the timeline and interpretative layers (that is, not connected to the rest of the encoding with *interpreted as* or any other arc). The benefit of this approach is that it allows for long but diffuse texts to be annotated as quickly as short, dense texts, where a "denser" text is one with more agentive implications per unit length. Detached from the highly laborious propositional modeling task, this variation in methodology allowed our annotators to create encodings for texts in a variety of genres including news, contemporary nonfiction, literary short fiction and epic poetry.

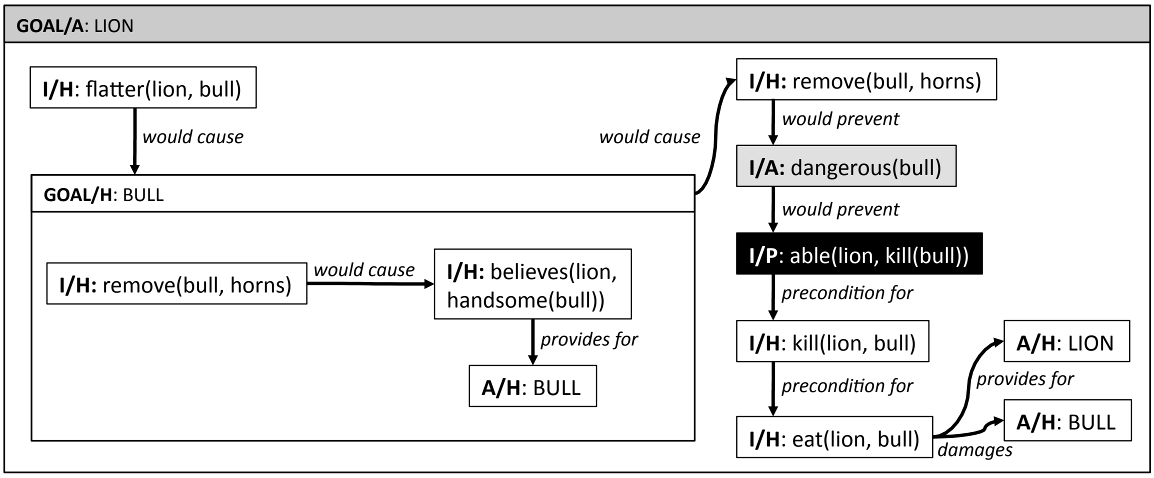

Figure 2: SIG encoding of the lion's plan at the moment of his scheming.

To assist in the annotation process, we extended the SCHEHERAZADE tool with an additional interface screen that presents a graphical canvas on which annotators can draw nodes and arcs. The system enforces the rules of the formalism and calculates actualization entailments, color-coding each node.

Our annotators consisted of six graduate students from our University's Department of English and Comparative Literature, as well as three undergraduate English majors. The training session took approximately 2-3 hours and consisted of an introduction to the SIG model, propositional modeling and the SCHEHERAZADE user interface. Annotators were given written guidelines and allowed to construct encodings at home, unsupervised and at their own pace. After finishing each encoding, they completed a survey in which they reported their satisfaction with encoding process for the particular text. The median time spent on an encoding was 1.25 hours (1 hour for Aesop fables, 2 hours for the longer texts).

Table 2 shows the DramaBank manifest, including the number of encodings elicited for each story within each collection. Every machine-readable encoding consists of a reproduction of its source text, the nodes and relations designated by the annotator, and the knowledge representations inside I and P nodes, if any. Table 2 also shows the average sizes of the encodings, in nodes and arcs per story, as well as the coverage—a clause is "covered" if the annotator designated a Text node to represent it and attached that node to a larger meaning structure in the timeline and interpretative layers.

Although the task is complex, annotators became comfortable with the tool after the training period. In the surveys, we asked them to report the tool's ease of use on a 5 point Likert scale, with 5 representing "easiest to use," and list specific aspects of the story that they were unable to encode to their satisfaction. For Collections B and C, the average usability score was 4.35. We also asked: "On a scale of 1 to 5, how satisfied are you that the system has encoded your interpretation of the story?" The average scores for Collections B and C were 4.26 and 4, respectively. Overall, the annotators reported satisfaction with the process, although the task was more intuitive for some annotators than for others (in a distant outlier, one annotator reported taking more than six hours to complete a single Collection B encoding). The annotation of the entirety of *Beowulf*, including 476 nodes covering more than 12,000 words (about 50%), was suggested and completed by an undergraduate specializing in medieval studies. She reported that the 15 hour project greatly enhanced her appreciation of the text by drawing out connections between disparate passages.

Short-form feedback by Collection B annotators indicated that propositional encoding was more difficult and constraining than agent-oriented annotation, with which they had little difficulty in comparison. This suggests that the set of agentive relations we have proposed are an accessible means for encoding the underlying intentional structure of a text. Conversely, among the annotators who worked with the longer Collection C texts, the most commonly reported issue was anxiety about engaging in the "mind reading" process and settling on a single interpretation. This is understandable, as several of these stories are ambiguous in terms of motivations of their characters, perhaps deliberately so. One annotator wrote that the system "forced [him] to choose an interpretation in a few places, where [he] might have wanted more room for ambiguity." To address this in future work, we will allow a single annotator to create plural readings of the same text, and indicate his or her confidence in each.

3.1. Evaluation

A potential downside to an expressive formalism such as the SIG is that low inter-annotator agreement can por-

Title	Encodings A	B	C	Nodes	Arcs	Coverage Words	%
Aesop (26 fables)	40	60		33.7	41.6	131	100.0
"An Alcoholic Case", F. Scott Fitzgerald			1	30	53	609	19.5
"Bahrain Protesters Say Security Forces Fire on Crowds", WSJ			1	35	53	246	21.4
The Battle of Maldon, anonymous			1	39	62	406	26.0
Beowulf, anonymous (trans. Slade (2011))			1	476	413	12,695	49.5
"The Gift of the Magi", O. Henry			1	46	61	422	20.3
"A Good Man Is Hard To Find", Flannery O'Connor			1	100	169	1,562	24.1
"The Lady with the Dog", Anton Chekhov			3	74.7	102.3	1,261	19.0
Sled Driver, Brian Shul (excerpt)			1	37	73	628	49.0

Table 2: Characteristics of DramaBank encodings.

tend challenges to machine annotation and useful aggregate analysis. We previously measured agreement in Collection A with a metric for the semantic similarity between modeled propositions (Elson and McKeown, 2010). For Collections B and C, the notion of inter-annotator agreement is complicated by the addition of the interpretative layer, which by its nature reflects a subjective take on the stated or unstated motivations driving the story's agents. The notion of a plural reading for an ambiguous text raises questions even about the notion of intra-annotator agreement. Differences between encodings are not only to be expected, but can themselves be a source of data about subjective differences in reception.

As an initial yardstick for agreement, though, we can at least expect that parallel encodings of the same story are more similar to one another than two encodings of different stories. As we described earlier, we defined *a priori* a set of 80 small reference encodings representing a non-exhaustive set of narrative scenarios, such as gain, loss, and the mixed blessing. We consider an encoding to cover one of these patterns if it contains a subgraph that is structurally isomorphic to it, i.e., if the pattern can map onto the encoding at least once. (We apply logical closure rules, such as transitive causality, to prevent minor variations from precluding valid mappings.) This allows us to measure inter-annotator agreement with respect to 80 particular features of narrative content. We found that homogeneous encoding pairs (different annotators, same source text) had a significantly higher cosine similarity between their feature vectors than heterogeneous encoding pairs (different source texts) to $p<.001$. If we apply Cohen's (1960) kappa statistic to the same data, considering each of the 80 features as a potential agreement or disagreement and taking the overall distribution of feature values as the basis for chance agreement, the result is $k=.55$. This represents moderate agreement.

As an alternate approach, we describe in a concurrent workshop paper (Elson, 2012a) a method for finding analogies between encoding pairs in a bottom-up fashion, without the use of hand-authored patterns. Based on the ACME model (Holyoak and Thagard, 1989) of analogy detection in connectionist representations of knowledge, it searches for the maximum common subgraph-isomorphism that maintains analogical consistency. The similarity between two encodings is then determined by the size of the largest isomorphic subgraph, normalized by story length, so that more similar pairs are those with greater structural overlap. Here, too, homogeneous encoding pairs in DramaBank are significantly more similar than heterogeneous pairs, to $p<.001$. While the specific points of agreement differ from pair to pair in this case, we have reason to believe that this interpretative-layer structural overlap indicates similarities between narratives that are meaningful—more so, in fact, than similarities found among propositional content in P and I nodes. In a separate evaluation, such structural features significantly outperformed our propositional similarity metric in fitting a linear regression model against a gold standard of the similarity of each story pair (as determined by separate raters comparing stories by reading the original texts). In other words, temporal and agentive SIG relations correlate better with ratings of story similarity than temporal and propositional annotations. These results encourage us to grow DramaBank's Collection C and pursue methods for automatic annotation and analysis of these relations.

4. Conclusion

We have proposed a novel set of discourse relations geared toward representing a theory-of-mind interpretation of narrative discourse. Our relations allow trained annotators to identify not only narrated events, with their temporal and modal relationships, but entities that give the stories a dramatic cohesion: agents with distinct beliefs, goals, plans, and affectual impacts (whether stated or implied by the text). Unlike prescriptive approaches to modeling narrative agents, the SIG is a descriptive model amenable to community corpus development and, ultimately, data-driven analysis of narrative corpora. While the theory of mind is not the only way to read and model a narrative text, this approach to discourse annotation lays the groundwork for new approaches to understanding narratives, their connections and their implications.

5. Acknowledgments

This material is based on research supported in part by the U.S. National Science Foundation (NSF) under IIS-0935360. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the NSF.

6. References

James F. Allen and George Ferguson. 1994. Actions and events in interval temporal logic. Technical Report 521, University of Rochester, Rochester, New York.

Mieke Bal. 1997. *Narratology: Introduction to the Theory of Narrative*. University of Toronto Press, Toronto, second edition.

Peer F. Bundgaard. 2007. The cognitive import of the narrative schema. *Semiotica*, 165(1–4):247–261.

Nathanael Chambers and Dan Jurafsky. 2008. Unsupervised learning of narrative event chains. In *Proceedings of the 46th Annual Meeting of the Association of Computational Linguistics (ACL-08)*, pages 789–797, Columbus, Ohio.

Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46.

David K. Elson and Kathleen R. McKeown. 2010. Building a bank of semantically encoded narratives. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC 2010)*, Malta.

David K. Elson. 2012a. Detecting story analogies from annotations of time, action and agency. In *Proceedings of the LREC 2012 Workshop on Computational Models of Narrative*, Istanbul, Turkey.

David K. Elson. 2012b. *Modeling Narrative Discourse*. Ph.D. thesis, Columbia University, New York City.

Amit Goyal, Ellen Riloff, and Hal Daumé III. 2010. Automatically producing plot unit representations for narrative text. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing (EMNLP 2010)*, Cambridge, Massachusetts.

Arthur C. Graesser, Murray Singer, and Tom Trabasso. 1994. Constructing inferences during narrative text comprehension. *Psychological Review*, 101(3):371–395.

Keith J. Holyoak and Paul Thagard. 1989. Analogical mapping by constraint satisfaction. *Cognitive Science*, 13:295–355.

V. S. Vernon Jones. 1912. *Aesop's Fables: A New Translation*. Avenel Books, New York.

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

William C Mann and Sandra A Thompson. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *Text*, 8(3):243–281.

Alan Palmer. 2007. Universal minds. *Semiotica*, 165(1–4):202–225.

Rashmi Prasad, Nikhil Dinesh, Alan Lee, Eleni Miltsakaki, Livio Robaldo, Aravind Joshi, and Bonnie Webber. 2008. The penn discourse treebank 2.0. In *Proceedings LREC 2008)*, Marrakech, Morocco.

James Pustejovsky, Patrick Hanks, Roser Saurí, Andrew See, David Day, Lisa Ferro, Robert Gaizauskas, Marcia Lazo, Dragomir Radev, Andrea Setzer, and Beth Sundheim. 2003. The timebank corpus. *Proceedings of Corpus Linguistics 2003*, pages 647–656.

Marie-Laure Ryan. 1991. *Possible Worlds, Artificial Intelligence and Narrative Theory*. Indiana University Press, Bloomington, Indiana.

Benjamin Slade. 2011. Beowulf on steorarume (beowulf in cyberspace). http://www.heorot.dk. Accessed 8/16/11.

Nancy L. Stein, Tom Trabasso, and Maria D. Liwag. 2000. A goal appraisal theory of emotional understanding: Implications for development and learning. In Michael Lewis and Jeannette M. Haviland-Jones, editors, *Handbook of emotions (2nd ed.)*, pages 436–457. Guilford Press, New York.

Soyoung Suh and Tom Trabasso. 1993. Inferences during reading: Converging evidence from discourse analysis, talk-aloud protocols, and recognition priming. *Journal of Memory and Language*, 32:279–300.

Paul van den Broek. 1988. The effects of causal relations and hierarchical position on the importance of story statements. *Journal of Memory and Language*, 27(1).

Rolf A. Zwaan and Gabriel A. Radvansky. 1998. Situation models in language comprehension and memory. *Psychological Bulletin*, 123(2):162–185.