Discovering Causal and Temporal Relations in Biomedical Texts

Rutu Mulkar-Mehta, Jerry R. Hobbs, Chun-Chi Liu and **Xianghong Jasmine Zhou**
University of Southern California
rutu@isi.edu, hobbs@isi.edu, jimliu@usc.edu, xjzhou@college.usc.edu

Abstract

In previous work on "Learning by Reading" we successfully extracted entities, states and events from technical natural language descriptions of processes. The research described here is aimed at the automatic discovery of causal and temporal ordering relations among states and events, specifically, among molecular and other events in biomedical articles. We have annotated causal and temporal relations in articles on the cell cycle, and we discuss our annotation guidelines and the issue of inter-annotator agreement. We then describe the natural language parsing and the inference system we use to extract these relations. We have created axioms manually for this system, focusing on temporal, causal and aspectual information and we have used semi-automatic means to augment these axioms. We have evaluated the performance of this system, and the results are promising.

Introduction

In three recent pilot projects on "Learning by Reading", sponsored by DARPA, we achieved some success in learning about complex processes and artifacts, and querying against this knowledge afterwards. Our systems were able to learn the components of chemical reactions, and the structure of simple events in the working of the heart and the workings of several kinds of engines. They did a good job of recognizing entities and simple events. But they did not recognize causal and temporal relations among events. Hence, our next project in the "Learning by Reading" area has been directed at the automatic discovery of these relations.

In our first pilot project (Mulkar, Hobbs, and Hovy 2007) we worked with chemistry textbook descriptions of chemical reactions and chemical entities. We were able to extract information about the structure of compounds before and after reactions from natural language. We used Powerloom (Chalupsky, Mac-Gregor, and Russ 2006) to evaluate the quality of information learned by the system. Our system was successful in answering "true/false" questions as well as a few basic "what" questions.

In the next pilot project (Mulkar et al. 2007; Barker et al. 2007) we adopted a new, more complex domain—descriptions in biological texts of the heart and how it works. The system was able to extract information about the parts of the heart and about the individual events each part was involved in. The backend knowledge base used for this system was Knowledge Machine (Clark, Harrison, and Thompson 2003). The system constructed complex models of the texts it read, with reasonable accuracy, and was able to answer successfully "what" and "how" questions.

In the third pilot project, we extended this work to several types of engines. It was again able to extract natural language information about entities, relations among entities, and simple events involving the entities, e.g., the flow of a fluid from one place to another. The backend reasoning engine used for this purpose was also Knowledge Machine. The final system was able to identify and create different models for different types of engines and correctly answer questions about their operation as described in the text.

All of these systems exhibited a limitation, however. They were able to learn information such as "blood is located in the heart" and "blood is located in different body parts". They were able to learn that "blood is oxygenated when it reaches the various body parts", and that "blood is not oxygenated when it returns to the heart". But they were unable to learn the temporal and causal relations between these events and states.

In our current research, described here, we are addressing the issue of identifying temporal and causal relations from text. Our aim is to learn the sort of information that would appear in a flow diagram of processes, where each state in the flow diagram is a snapshot of the activity at a particular instant of time, and the temporal sequence is represented by the state transitions. For a domain, we have chosen biomedical articles, because they are typically very rich in causal and temporal relations, often with several such relations in every sentence.

Recognizing Causal and Temporal Relations: Current Technology

In the natural language processing community, there have been a number of studies on temporal relations. For example, (Filatova and Hovy 2001) find references to specific dates and times and propagate those to subsequent clauses by interpreting deictic and other temporal expressions. However, their experiments have been on news corpora, and incorporate the timestamps in the article. Such information is unavailable or irrelevant in most scientific literature, so it is not applicable in our application.

(Lapata and Lascarides 2004) have examined sentence-internal temporal relations. However, their work focuses on sentences containing a main and subordinate clause. They attempt to discover the temporal relation between the two events, rather than looking at the often intricate temporal relations within single clauses.

(Bethard et al. 2008) have recently published their work on building a corpus containing temporal and causal relations. However, they have limited their annotations to events in sentences conjoined by "and", and have worked only with newspaper articles.

(Mani et al. 2006a) and later (Chambers, Wang, and Jurafsky 2007) have also worked on finding temporal relations in newspaper articles using both handcoded rules and machine learning techniques. They used aspectual and tense information for analyzing the temporal relations. The temporal ordering toolkit, TARSQI (Temporal Awareness and Reasoning Systems for Question Interpretation) (Pustejovsky, Mani, and Hobbs 2007; Verhagen 2004), is the only complete toolkit that is freely available for use. We applied this toolkit for the biomedical domain on a set of 40 sentences, and obtained a good score for non-biological event recognition (Recall = 70.6%, Precision = 97.8%) , but a low score for non-biological event ordering (Recall = 47.0%, Precision = 47.6%), and an even lower score if biological events are included (Recall = 26.52%, Precision = 27.03%)

TARSQI has several shortcomings with respect to scientific texts. It often determines event order by the succession of tenses. While this information is pivotal in a narrative news text, scientific texts rarely convey temporal relations with their usage of tenses. TARSQI also uses text order to predict temporal relations, but in scientific texts, event order and text order are often at odds. In the sentence

*The **progression** into G2 phase depends mainly on another wave of cyclin **production**.*

TARSQI determines the temporal relation to be "progression → production". However, it is clear to a reader that the progression is dependent on cyclin production, and hence the production must happen before the progression.

Causal and Temporal Information in Biomedical Text

Textual data sometimes contains overt temporal markers that convey when a temporal event takes place. For example, in the sentence,

*XPD appears to be degraded in wild-type embryos **between** prophase and metaphase of the first cell division **after** the onset of zygotic gene expression, which **coincides** with a redistribution of CDK7 from the cytoplasm to the nucleus*,

the presence of temporal markers "after","between" and "coincides" helps us pin down when XPD appears to be degraded. But there are other ways temporal information is conveyed in biomedical text.

Texts conveying causality: Causal words play a major role in determining the order of occurrence of events. If event A causes event B, event A occurs before event B. For example, In the sentence

*In Drosophila, inactivation of the CDK7 mutant **causes** embryonic and larval lethality and a block to mitosis in the germline.*

we know that "larval lethality" and "block to mitosis" occur after the "inactivation of the mutant". Another example is as follows:

*Mutations in the gene encoding XPD **lead to** dysregulation of nuclear-receptor-dependent gene expression.*

Here "lead to" is the causal marker indicating that "mutations in a specific gene", occur before "dysregulation".

The following is an example sentence that is extremely rich in causal/temporal relations but has almost no overt temporal markers. The causal markers are in **bold**.

*In budding yeast, **when** the essential function of CAK is bypassed **by** mutations within the CDC28 gene encoding CDK1, cak1 mutants display defects during vegetative growth **as a result of** diminished function of other CAK targets involved in transcription, such as Bur1 and Ctk1, and meiotic defects associated with impaired activation of the sporulation-specific kinase Ime2.*

This sentence is rich in temporal information, but contains only one explicit temporal marker ("during"). However, based on the causal relations in the text, the following temporal ordering can be derived. (Here → represents temporal order and == represents simultaneous occurrence)

mutations within the CDC28 gene → function of CAK is bypassed → diminished function of other CAK targets AND meiotic defects → cak1 mutants display defects == vegetative growth

Aspectual words: These words represent an aspect of a higher level event. For example, phrases like "entry into", "exit from" represent an aspect of the main event from which an entry or exit might take place. In the sentence

In S. pombe, the Wee1 kinase allows ***entry*** *into mitosis when cells reach a minimum threshold size.*

we are given the information that "reaching a minimum threshold size" is a precondition for beginning the process of "mitosis" but does not guarantee the entire process of mitosis will take place. Thus we get the temporal relation stating that "once a minimum threshold is reached", "mitosis will begin", but the entire process of mitosis might not occur.

Enabling words or words conveying dependence: If an event B is dependent on event A, event B does not happen if event A does not happen, so event A occurs before event B. Phrases conveying dependence are "A critical for B", "A necessary for B", "A required for B". (Sample axioms in Table 1.) An example is:

Regulation of cell cycle progression is another general stress response ***critical for*** *cell survival.*

(Hobbs 2005) describes these sets of words to be a part of the causal complex of an action. For example, proper working of a power plant is a necessary condition, but not a direct cause for a light to turn on. The flipping of a switch is the direct cause for the light to turn on.

Words conveying control: If event A controls event B, A is a necessary condition for B, and A will not occur after B. Examples of this type are "A governs B", "A is responsible for B", "A promotes B", "A produces B". (Sample axioms are given in Table 1.)
Two examples of this case are the following:

Dephosphorylation of these residues by the Cdc25 phosphatase is the key event ***governing*** *the initiation of mitosis.*

Activated Hog1 ***promotes*** *the delocalization of Hsl7 from the neck that results in Swe1 accumulation.*

Annotation Effort

The different types of temporal information in text, described in the previous section, can be classified into five broad temporal categories, represented in Figure 1. These categories are aligned with the relations used by (Mani et al. 2006b) and TARSQI (Pustejovsky, Mani, and Hobbs 2007).

Annotation Guidelines

Event Identification: As previously mentioned, TARSQI (Pustejovsky, Mani, and Hobbs 2007) has a high precision and moderate recall for event recognition. As a result, it was used for initial event recognition. It, however, did not mark any biological events; these were manually annotated.

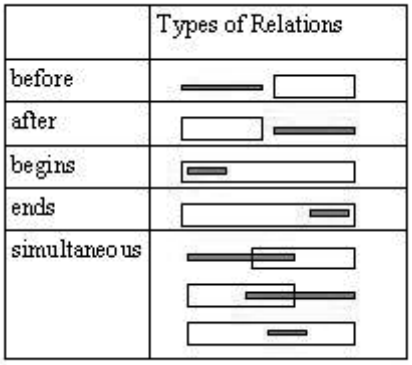

Figure 1: Temporal relations identified in text.

What do we want to annotate: Our aim was to annotate the order of occurrence of events or a set of events within given sentences. We did not annotate relations between events across sentences.
Annotation Tool: RSTTool (Version 3.0) by Michael O'Donnell (http:// www.wagsoft.com /RSTTool/) was used for annotation. This tool aids in text segmentation, text structuring, relation specification, and maintenance of statistics for relations.
How to Annotate: Given a sentence, each annotator first identified the individual sentence fragments that convey a temporal relation. The annotator then created a temporal link connecting the head verb or event noun of one fragment to the head of the other.

Inter-Annotator Agreement

A set of 40 sentences were randomly selected from the test corpus. These were annotated by two annotators after reading the annotation guidelines. There were a total of 215 events in the sentences. Annotator A had a biology background, and annotated 90 temporal relations. Annotator B had a computer science background and annotated 73 relations.

The agreement between the annotators was computed by taking annotator A's relations as the gold standard. B's annotation was considered correct if A explicitly annotated the relation or it could be inferred from A's annotations (e.g., by transitivity). B was deemed to have detected a relation annotated by A if it could be inferred from B's annotations. The chief reasons for discrepancies between the annotators were that A sometimes included relations from domain knowledge that the text did not support, and that A sometimes overlooked relations involving non-biological events.
Annotator B achieved a precision of **69.8%** and recall of **53.3%**

The Kappa statistic (Krippendorf 1980) for inter-annotator agreement in this task with respect A's annotations is 0.95, but this is not particularly illuminating, since the space of events is overwhelmingly dominated by event pairs for which no temporal relation holds.
A combined gold standard was created by reconciling the two sets of annotations through discussion and producing the integrated annotations. This gold standard was used for subsequent evaluations.

Baseline: To compute chance agreement, event pairs were randomly selected from the text (each pair occurring within the same sentence) and a random relation was picked for each of these event pairs. The number of event pairs selected was based on the event-to-relations ratio in the gold standard; the ratio was 2.34, which meant that there should be one relation roughly for every two events. In a sense this baseline is as if an annotator knew roughly how many relations to produce, but did not read the text. Ten sets of such annotations were randomly created, and their average precision and recall compared with the gold standard was **4.43%** and **6.29%** respectively.
In a related experiment, the baseline generation system was given all the events from the gold standard that were involved in a temporal relation. The system picked two events out of these known events and randomly generated a temporal relation between them. Though this method relied on too much information for generating a chance agreement, the average precision and recall still remained low - **11.64%** and **11.0%** respectively.
Finally, a third technique for baseline generation was implemented, where the system was provided with all the events involved in a temporal relation. Event pairs were randomly selected within each sentence and temporal relations were generated between them based on the order in which the events were mentioned in the sentence. The average precision and recall in this case were again low: **10.30%** and **8.33%**. The low baselines give some idea about the complexity of the problem.
Figure 2 shows the gold standard, the baseline, and the annotations by the 2 annotators.

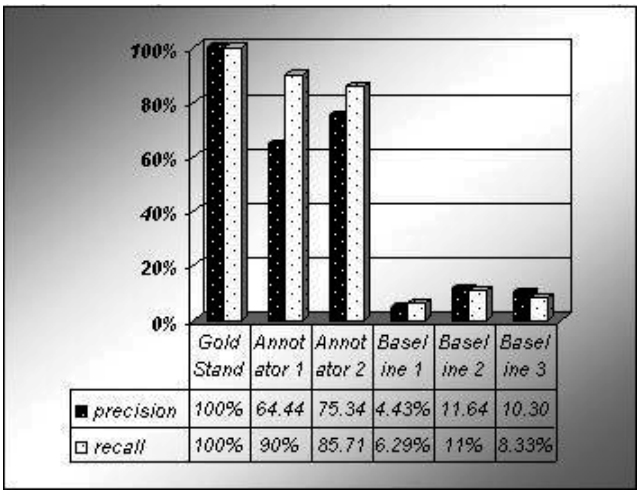

	Gold Stand	Annotator 1	Annotator 2	Baseline 1	Baseline 2	Baseline 3
precision	100%	64.44	75.34	4.43%	11.64	10.30
recall	100%	90%	85.71	6.29%	11%	8.33%

Figure 2: Gold Standard, Baseline and Annotations

The Natural Language Pipeline

Our ultimate goal is to identify the temporal and causal relations automatically. Our natural language pipeline consists of three steps: parsing the text, generation of a shallow logical form, and deriving temporal inferences from this logical form by abduction.

Parsing and Logical Form Generation

The Charniak parser (Charniak 1999) was used to parse the text. All references to tables, figures and citations were removed from the text before parsing it. No other simplification or changes were made to the text.

This parse is first converted to an XML format, with lexical features inserted from a large dictionary. The dictionary was a locally developed dictionary of 400,000 words derived from SRI's DIALOGIC system, COMLEX, Maurice Gross's dictionary, and some local modifications and additions. The parse is then converted to a binary tree in which logical and structural information is propagated to the appropriate nodes. Rules in LFToolkit (Rathod and Hobbs 2005) then convert this to a shallow logical form (Hobbs 1985) making explicit the predicate-argument relations implicit in the syntax. Below is the shallow logical form for the sentence: *Exposure of cells to stress results in rapid activation of MAPKs*.

exposure-nn(x2) & of-in(x2,x10) & cell-nn(x10) & to-in(x2,x6) & stress-nn(x6) & result-vb'(e0,x2) & in-in(e0,x8) & rapid-adj(x8) & activation-nn(x8) & of-in(x8,x12) & mapks-nn(x12)

Here -nn,-vb,-in tags appended at the end of each word indicate the part of speech of the word (noun, verb and preposition, respectively). E.g. The variable *x2* represents the "exposure" and *e0* represents the "result" eventuality. The exposure *x2* is the subject of the "result" eventuality *e0*. The object of *e0* is the activation *x8*.

Inferring of Temporal Relations

An abductive inference engine, Mini-TACITUS (Hobbs et al. 1993; Mulkar, Hobbs, and Hovy 2007) was used for inferring temporal relations in text. It attempts to find the best possible explanation for the sentence, given a set of axioms as the knowledge base. In this case the axioms used were general commonsense knowledge conveying the relations among causality, aspect and temporal order.
Example: The above sentence contains the subphrase "results in".

The following axiom relates the verb "result" to the underlying predicates CAUSE and BEFORE:

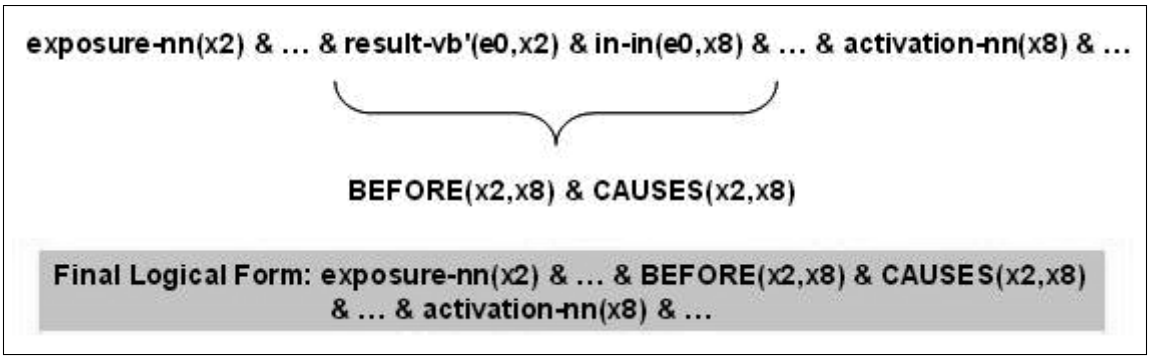

Figure 3: Example of abductive process

CAUSES(x3,x4) & BEFORE(x3,x4) → result-vb'(e0,x3) & in-in(e0,x4)

This says that if *x3* is before *x4* and causes *x4*, then *x3* results in *x4*. The causality is an explanation for the explicit "result" relation. The axiom brings along the BEFORE relation as a side effect in abductive inference. This is illustrated in Figure 3.

Axiom Creation Process

Method 1: Manual

A total of 66 axioms were manually created to handle the causal and aspectual relations in 190 sentences from two biomedical articles on the cell cycle. A sample of this set of axioms is shown in Table 1. Table 2 shows the results of using these axioms on a test set of 40 sentences.

The recall of 29.4% is low, but this is an indication that there are more ways to encode temporal information than the causal and aspectual rules we have authored so far. When we look only at temporal relations that are signaled causally, the recall goes up to 64.3%. When we restrict ourselves to those sentences for which the Charniak parser got the correct parse (column NBP for "No Bad Parses"), the recall was 90.0%. The Charniak parser was not trained on biomedical literature, so there is a clear path to improvement there.

The precision of 67.7% is not especially high, but almost all of the errors can be attributed to bad parses. When these are eliminated, the precision is 100% in both the causally and aspectually encoded events.

The recall for the aspectual class is low, but that is the result of having too few aspectuals in the test data.

Method 2: Semi-Automatic

To augment our set of axioms beyond what was required for the 190 sentences used in the manual axiom creation, we employed automated pattern extraction on a large dataset of over 2 billion words. A set of 12 seed words (49 total words including verb variations) containing causal relations and 15 seeds (31 including verb variations) containing aspectual relations were manually selected, and string patterns of the form X SEED Y were extracted from three different releases of the LDC corpus. The algorithm from (Bhagat and Ravichandran 2008) was implemented with the standard settings from the paper. It was determined what X's and Y's occur with the seed and then what other words and phrases occur in the same contexts. For example, with the seed pattern "X results in Y", we retrieved the string "X triggers Y". We would expect this to be a high recall, though low precision, way of determining how causality is expressed in text. It should be mentioned that our corpus is 10% the size of the original corpus, and we have not made an effort to tune the parameters for the corpus size or the task. The results produced by this implementation should be interpreted in this light.

Results for causal relations: The system returned a total of 308 phrases, of which 182 were syntactic variations of the original seed relations. This result by itself is encouraging because this implies that the coverage with our manually constructed axioms after 2 articles is better than 59%. The remaining 126 patterns could be collapsed into 100 pattern classes, where each class contained syntactic variations of the same root word (e.g. "response to", "responds to" and "will respond to" all lie in the same class). 79 of these pattern classes were irrelevant (e.g. "X is illegal in Y" , "X commemorating Y", "X were named directors of this Y"). Only 9 new valid pattern classes were discovered in this experiment - "allow", "back", "base","devote", "encourage", "prompt", "rely", "rule" and "trigger". All of these were of relatively low frequency.

Results for aspectual relations: A total of 320 phrases were returned. These fell into 106 pattern classes, of which 79 were completely irrelevant (e.g. "marched in", "four-game winning"), 15 classes were variations of the seeds originally supplied, and 12 new patterns were discovered (e.g. "conclude", "cease","erupt","kick off").

The results of this experiment are encouraging, because they indicate that not very many new patterns occurred in a 2 billion word corpus, over and above what we had found manually in only about 5500 words of text. Moreover,

	Causal Axioms
1	CAUSES(x3,e4) & BEFORE(x3,e4) → in-in(e4,x2) & response-nn(x2) & to-in(x2,x3)
2	CAUSES(x0,x8) & BEFORE(x0,x8) → result-vb'(e0,x0) & in-in(e0,x8)
3	CAUSES(x1,x3) & BEFORE(x1,x3) → lead-vb'(e1,x1) & to-in(e1,x3)
4	CAUSES(e4,x3) & BEFORE(x1,x3) → in-in(e4,x2) & order-nn(x2) & to-in(x2,x3)
5	CAUSES(x15,e14) & BEFORE(x15,e14) → upon-in(e14,x15)
	Aspectual Axioms
6	BEGINS(e1,e2) → start-vb'(e1,x1,e2)
7	BEGINS(x1,x2) → progression-nn(x1) & into-in(x1,x2)
8	BEGINS(x1,x2) → onset-nn(x1) & of-in(x1,x2)

Table 1: Sample axioms for Mini-TACITUS. It is true that CAUSE implies BEFORE, but in an abductive framework such an axiom would function primarily to interpret instances of a "BEFORE" relation as causal. Since we do not forward chain, the "BEFORE" relation would not be inferred once we assumed "CAUSE". On the other hand, placing an extra conjunct in the antecedent is one way to capture a (forward-chaining) side effect of an assumption in weighted abduction.

	All	**Causal**		**Aspect**	
	All	**All**	**NBP**	**All**	**NBP**
Pre.	67.7%	69.2%	**100%**	33.3%	**100%**
Rec.	29.4%	64.3%	**90.0%**	12.5%	16.7%

Table 2: Results of axiom implementation on a test set

it suggest that while simple automatic techniques for the discovery of indicators of temporal relations are useful for boosting recall, they require manual filtering to eliminate substantial amounts of irrelevant material.

Completeness

In order to determine when the manual construction of axioms is likely to converge, we manually created the axioms incrementally for ten sentences at a time. Figure 4 shows the number of axioms that had been created after each set of ten sentences. There were a total of 66 axioms created at the end of 190 sentences.

The first fifty sentences resulted in 38 axioms, while the last 50 sentences resulted in 7 new axioms This is a significant drop-off. This of course is not a proof of convergence but it is suggestive.

In addition, the results of the semi-automatic method of axiom construction (Method 2) are strongly suggestive of convergence, because they yielded very few axioms over and above our manual efforts. It is true that the paraphrasing algorithm we used yielded low precision, but it did promote high recall, which was our goal.

Conclusions and Future Work

There are a number of obvious places where the work described here can be augmented or improved. We would like to refine our annotation guidelines and software to get a higher level of inter-annotator agreement. In the construction of axioms, there are other classes, besides cause and aspect, of indicators of temporal information, and we would like to extend our efforts to these. Our studies

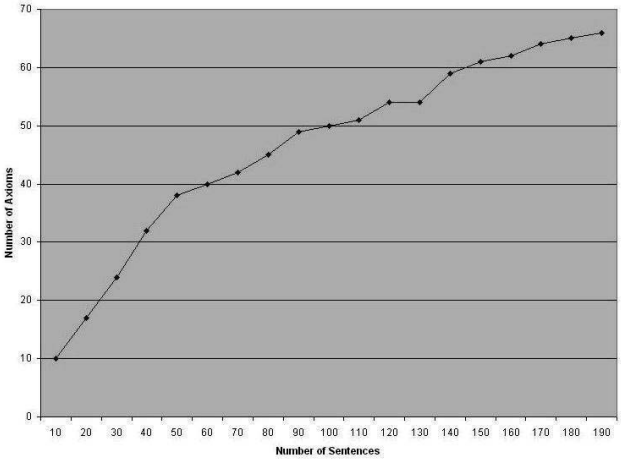

Figure 4: Axioms generated every 10 sentences

suggest convergence in the sets of causal and aspectual axioms, but we would like to continue the process to the point where we actually see convergence.

Learning about processes from reading natural language texts describing the processes requires not only recognizing entities, states and events, but also requires temporal and causal relations among these states and events. We believe the research described here represents a significant step in that direction.

Acknowledgments

We have profited from discussions with Rahul Bhagat and Eduard Hovy in this research. It has been partially funded by NSF grant 0747475.

References

Barker, K.; Agashe, B.; Chaw, S. Y.; Fan, J.; Friedland, N. S.; Glass, M.; Hobbs, J. R.; Hovy, E. H.; Israel, D. J.; Kim, D. S.; Mulkar-Mehta, R.; Patwardhan, S.; Porter, B. W.; Tecuci, D.; and Yeh, P. Z. 2007. Learning by read-

ing: A prototype system, performance baseline and lessons learned. In *AAAI*, 280–286. AAAI Press.

Bethard, S.; Corvey, W.; Klingenstein, S.; and Martin, J. 2008. Building a corpus of temporal-causal structure. *In Language Resources and Evaluation Conference (LREC)*.

Bhagat, R., and Ravichandran, D. 2008. Large scale acquisition of paraphrases for learning surface patterns. *Proceedings of ACL'08*.

Chalupsky, H.; Mac-Gregor, R. M.; and Russ, T. A. 2006. Powerloom manual. *In http://www. isi. edu/ isd/LOOM/ Power- Loom/ documentation /manual.pdf*.

Chambers, N.; Wang, S.; and Jurafsky, D. 2007. Classifying temporal relations between events. In *ACL*. The Association for Computer Linguistics.

Charniak, E. 1999. A maximum-entropy-inspired parser. Technical report, Providence, RI, USA.

Clark, P.; Harrison, P.; and Thompson, J. 2003. A knowledge-driven approach to text meaning processing. In *In Proceedings of the HLT Workshop on Text Meaning Processing*, 1–6. ACL Press.

Filatova, E., and Hovy, E. 2001. Assigning time-stamps to event-clauses. *Proceedings of the workshop on Temporal and spatial information processing* 1–8.

Hobbs, J. R.; Stickel, M.; Martin, P.; and Edwards, D. D. 1993. Interpretation as abduction. *Artificial Intelligence Vol. 63, Nos. 1-2, pp. 69-142, 1993* 95–103.

Hobbs, J. R. 1985. Ontological promiscuity. *Proceedings of the 23rd annual meeting on Association for Computational Linguistics* 60–69.

Hobbs, J. R. 2005. Toward a useful concept of causality for lexical semantics. *Journal of Semantics* 22:181–209(29).

Krippendorf, K. 1980. Content analysis: An introduction to its methodology. *Sage Publications*.

Lapata, M., and Lascarides, A. 2004. Inferring sentence-internal temporal relations. *In Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics, Boston, MA*.

Mani, I.; Verhagen, M.; Wellner, B.; Lee, C. M.; and Pustejovsky, J. 2006a. Machine learning of temporal relations. *ACL '06: Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the ACL* 753–760.

Mani, I.; Verhagen, M.; Wellner, B.; Lee, C. M.; and Pustejovsky, J. 2006b. Machine learning of temporal relations. In *ACL-44: Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, 753–760. Morristown, NJ, USA: Association for Computational Linguistics.

Mulkar, R.; Hobbs, J.; Hovy, E.; Chalupsky, H.; and Lin, C.-Y. 2007. Learning by reading: Two experiments. *Proceedings of the IJCAI 2007 workshop on Knowledge and Reasoning for Answering Questions*.

Mulkar, R.; Hobbs, J.; and Hovy, E. 2007. Learning from reading syntactically complex biology texts. *Proceedings of the AAAI Spring Symposium Commonsense'07. Stanford University, CA*.

Pustejovsky, J.; Mani, I.; and Hobbs, J. 2007. Temporal awareness and reasoning systems for question interpretation. *http://www.timeml.org/site/tarsqi/*.

Rathod, N., and Hobbs, J. 2005. Lftoolkit. *http://www.isi.edu/~hobbs/LFToolkit/index.html*.

Verhagen, M. 2004. Times between the lines. *Ph.D. thesis. Brandeis University. Waltham, MA, USA*.