

Machine Reading through Textual and Knowledge Entailment

Andrew Hickl

Language Computer Corporation
1701 North Collins Boulevard
Richardson, Texas 75080
andy@languagecomputer.com

Sanda Harabagiu

Human Language Technology Research Institute
University of Texas at Dallas
Richardson, Texas 75080
sanda@hlt.utdallas.edu

Abstract

While information extraction and textual entailment systems have greatly enhanced the amount of knowledge available from a text corpus, the next generation of natural language understanding systems – such as Machine Reading systems – will need to employ dedicated mechanisms to ensure that acquired knowledge is *consistent* with the previous commitments encoded in the system’s knowledge base. This paper introduces a new technique for knowledge acquisition for Machine Reading systems, known as *Reading by Entailment* (RbE), which uses two complementary systems for recognizing entailment relationships in order to identify the set of hypotheses that can be inferred from a text corpus. Entailment hypotheses are generated by a random walk model that operates on texts and informs a plausible model of knowledge entailment. Knowledge entailment is cast as a maximal knowledge relevance problem which operates on knowledge mappings.

1. Introduction

The recent attention paid to the task of recognizing textual entailment (RTE) (Dagan *et al.* 2006; Bar-Haim *et al.* 2006) has led to the development of a number of systems capable of accurately recognizing textual entailment (TE) relationships in natural language texts (Hickl *et al.* 2006; Tatu *et al.* 2006). While RTE systems do not provide the same robust natural language understanding capabilities sought by Machine Reading systems, we argue that systems for recognizing textual entailment can enhance the amount of knowledge available from a text collection by identifying text passages whose content supports the inference of one or more textual hypotheses.

As currently defined, the RTE task has focused exclusively on entailment based solely on the forms of lexicosemantic and pragmatic information that is computable from a text, without explicitly leveraging structured sources of domain or world knowledge. Many current approaches to the RTE task, however, have looked to incorporate new resources that expand the amount of textual knowledge available to a system. In the 2006 PASCAL Recognizing Textual Entailment Challenge (Bar-Haim *et al.* 2006), sys-

tems successfully exploited a number of different sources of textual knowledge, including (1) wide-coverage lexical resources such as WordNet or FrameNet (Delmonte *et al.* 2006), (2) databases of paraphrases (Hickl *et al.* 2006; de Marneffe *et al.* 2006), (3) libraries of natural language axioms (Tatu *et al.* 2006), or (4) collections of additional examples of textual entailment that were used to train statistical classifiers (e.g. (Hickl *et al.* 2006)) or to build formal models for RTE (Bos & Makert 2006).

While RTE systems have focused on acquiring forms of knowledge from text, the next generation of Machine Reading systems will need to employ dedicated mechanisms to ensure that acquired knowledge is *consistent* with the previous commitments encoded in the system’s knowledge base (KB). We assume that knowledge is consistent with regards to a particular model iff the truth of the proposition can be reasonably inferred from the other knowledge commitments of the model. Although we believe that approaches based on model checking and theorem proving hold much promise for formally checking the consistency of acquired knowledge, we believe that the recognition of consistent knowledge could be modeled using the many of the same types of approximation techniques that have been successfully employed for the task of recognizing textual entailment.

In this paper, we argue that frameworks for RTE could be repurposed in order to recognize entailment relationships between commitments stored in a KB and any of a set of hypotheses generated from text. We introduce a relationship, known as *knowledge entailment*, that can hold between a commitment derived from an existing knowledge base and a natural language hypothesis generated from a text corpus. As with (Dagan *et al.* 2006)’s definition of textual entailment, we propose that a hypothesis is considered to be *knowledge entailed* iff there exists some knowledge commitment available in the KB could be used to lead a system to infer the truth of the hypothesis for all of the models derivable from the KB. As with TE, we expect that the lack of knowledge entailment between a commitment and a hypothesis does not imply that the content of the hypothesis is inconsistent with the content of the KB, but rather that commitment in question does not encode the knowledge necessary to support the entailment of the hypothesis.

We believe that systems for recognizing textual entailment and recognizing knowledge entailment (RKE) can be

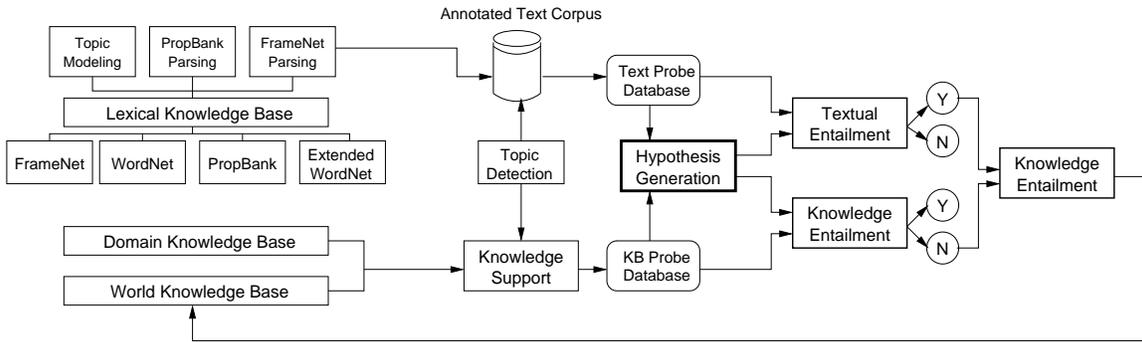


Figure 1: Architecture of a *Reading by Entailment* System.

coupled with systems for generating hypotheses from text in order to provide a robust framework for acquiring the forms of knowledge needed by a Machine Reading system. We introduce a new technique for knowledge acquisition, known as *Reading by Entailment* (RbE) – which uses the output of systems for recognizing textual and knowledge entailment in order to identify the set of hypotheses that can be inferred from a text corpus, regardless of the amount of previous knowledge available to the system.

2. Reading by Entailment

In order to perform deep semantic understanding of a corpus, Machine Reading applications need to employ combinations unsupervised techniques for capable of acquiring diverse forms of knowledge from a text. In this section, we describe a new framework that we have developed – known as Reading by Entailment (RbE) – which uses the recognition of entailment relationships in order to acquire and consolidate knowledge from a text collection for a Machine Reading application.

Our RbE system operates over text collections that have been richly annotated with a wide variety of lexico-semantic information, including (1) semantic dependency information derived from semantic parsers based on PropBank annotations, (2) semantic frame information made available from FrameNet parsers, and (3) automatically generated representations of the topics encoded in a document collection, such as those proposed by (Lin and Hovy 2000) or (Harabagiu 2004). Once annotation is complete, the RbE system divides the available text collection into a set of sentence-length *text probes* which can be used to evaluate hypotheses sent to a Textual Entailment module. In a similar fashion, commitments from the system’s KB are also selected as *knowledge probes* and sent to a Knowledge Entailment module to evaluate hypotheses as well.

The RbE system then generates a single set of hypotheses from the text corpus which can be evaluated against the set of extracted text and knowledge probes. Recognizing knowledge can be performed by a mechanism that generates hypotheses from each machine-read text. In order to generate hypotheses, we have created a model inspired by the method for decomposing complex questions introduced in (Harabagiu *et al.* 2006). In that work, the decomposition of complex questions was cast as a Markov Chain which performs a random walk over a bipartite graph of relations which have been identified in the text collection. In a sim-

ilar fashion, we cast the process of hypothesis generation from a text as another Markov Chain that performs a random walk over a bipartite graph of knowledge operators and conceptual representations that constitute the coercions from a text. (Full details of our method for hypothesis generation are provided in Section 4.) Each hypothesis is then paired with each of the text and knowledge probes and are then considered by systems for recognizing textual and knowledge based entailment.

Following (Hickl *et al.* 2006), we perform RTE using a classifier that estimates the likelihood that a hypothesis is textually entailed by a text probe. (Details of our system for RTE are provided in Section 3.) We currently perform the recognition of knowledge entailment (RKE) using an extension of the probabilistic abduction model introduced in (Raina *et al.* 2005). In this work, we have sought to identify two types of knowledge mapping functions which have allowed us to estimate the cost of performing an abduction, given either (1) the current state of the KB or (2) the state of the KB after the incorporating new knowledge associated with the hypothesis. (Details of our system for RKE are provided in Section 5.)

When a hypothesis is textually entailed by a text probe, the RbE system infers that the knowledge available in the hypothesis is textually supported and therefore merits consideration for inclusion in the knowledge base. However, before the knowledge associated with the hypothesis can be added to the KB, RbE must determine whether knowledge encoded by the hypothesis is already included in the KB or is somehow contradicted by the information stored in the KB. In order to perform this validation, textually-entailed hypotheses are therefore first sent to a *Knowledge Consolidation* module before they are added to the KB.

When a hypothesis is knowledge entailed by a probe, however, we expect the RbE system can reasonably infer that the knowledge associated with the hypothesis is already contained in the KB and does not represent a novel contribution to the knowledge available to the system. With hypotheses that fail to be entailed by any of the selected knowledge probes, we expect that the hypothesis represents information that is either (1) previously unattested in the KB or (2) is at odds with the commitments of the KB. These hypotheses are sent to *Knowledge Consolidation* for further evaluation.

We perform *Knowledge Consolidation* in two ways. First, hypotheses that are TE by at least one probe are then re-evaluated against the set of knowledge probes. If the hy-

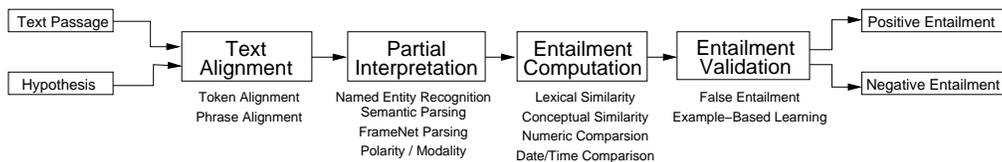


Figure 2: Schematic Architecture of a Textual Entailment System.

pothesis is KE by at least one probe, the system infers that the knowledge associated with the hypothesis is already subsumed by the KB. If the hypothesis does not pass knowledge entailment for any of the probes, it is then sent to a system for recognizing contradiction (modeled after the work of (Harabagiu *et al.* 2006) on recognizing textual contradictions) in order to determine whether it is directly contradicted by any of the previous commitments of the KB. If it is not deemed to be inconsistent with the previous commitments of the KB, the hypothesis is then added to the knowledge base as a new commitment. Second, once a set of textually entailed hypotheses have been added to the KB, the system then re-evaluates the hypotheses that were not previously knowledge entailed against the set of knowledge probes plus the set of new commitments added to the knowledge base; if the hypothesis still fails to be knowledge entailed by any of the new set of probes being considered, the hypothesis is added to the knowledge base as a new commitment.

3. Recognizing Textual Entailment

Much recent work has demonstrated the viability of supervised machine learning-based approaches to the acquisition of robust forms of textual inference such as textual entailment or textual contradiction. In these systems, textual inference is recognized using a classifier which evaluates the probability that a particular inferential relationship exists between two text passages using models based on a variety of statistical, lexico-semantic, or structural features.

In the RTE task, systems determine whether the meaning of a sentence (referred to as a *hypothesis*) can be inferred from the meaning of a second, often longer, text passage (known as a *text*). Two representative examples from the 2006 PASCAL RTE Challenge are provided in Table 1.

Judgment	Example
YES	Text: Scientists have discovered that drinking tea protects against heart disease by improving the function of the artery walls. Hypothesis: Tea protects from some diseases.
NO	Text: Aspirin, an inexpensive drug, helps protect survivors of heart attack and stroke from subsequent heart attacks and death. Hypothesis: People experienced adverse effects while taking aspirin.

Table 1: Examples from the 2006 PASCAL RTE Challenge.

In the first example in Table 1, RTE systems must recognize that one of the propositions encoded by the text – *tea protects against heart disease* – is a true statement, and one whose acceptance should lead a user to uncontroversially accept the truth of the proposition encoded by the hypothesis, *tea protects from some diseases*. In contrast, in the second example, part of the implication of the text – namely, that *as-*

pirin protects people – is directly contradicted by the meaning of the hypothesis, which can be reasonably paraphrased as *aspirin causes people to experience adverse effects*.

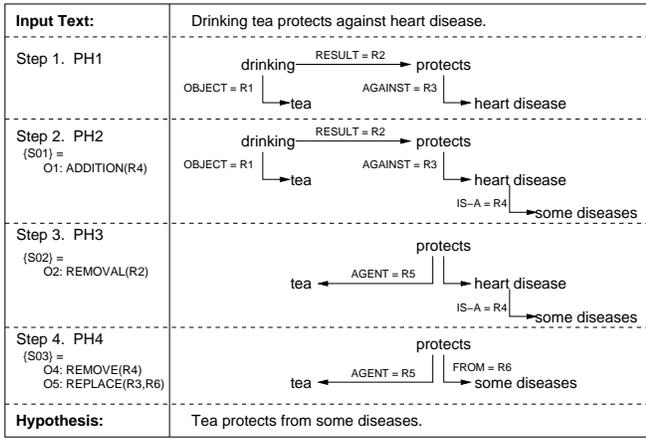
While the interpretation of these two examples would appear to require a substantial natural language understanding capacity, we believe that statistical approach to RTE that leverages a rich substrate of lexico-semantic annotations can be used to reliably estimate the probability of a textual entailment relationship. In previous work (Hickl *et al.* 2006), we have adopted a four-stage approach to the recognition of textual entailment (TE) relationships in text. (The schematic architecture of our system is depicted in Figure 2.)

In our system, pairs of texts and hypotheses are initially submitted to a *Textual Alignment* module, which utilizes a classification-based approach in order to identify pairs of constituents which convey similar semantic information. For each pair of tokens (or phrases) extracted from the text (w_t) and the hypothesis (w_h), the alignment classifier estimates the probability that an alignment relationship can be established between w_t and w_h . After each pair of constituents is classified, we use an incremental beam search in order to identify the set of token- or phrase-level alignments which maximizes the sum of the alignment probabilities for a text-hypothesis pair. Information from Textual Alignment is also used in order to capture portions from each of a pair of texts that could be related via one or more phrase-level alternations or “paraphrases”. The most likely pairs of aligned constituents identified by this module are then sent to a separate *Paraphrase Acquisition* system, which extracts sentences that contain instances of the aligned constituents from the WWW.

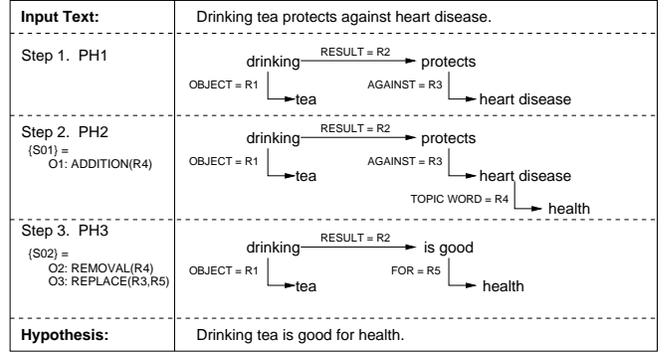
Following alignment and paraphrase acquisition, texts and hypotheses are annotated with a wide variety of semantic and pragmatic information by a *Partial Interpretation* module. As described in (Hickl *et al.* 2006), this module labels named entities, normalizes temporal and spatial expressions, resolves instances of coreference, and annotates predicates with polarity, tense, and modality information. In addition, we use semantic parsers based on PropBank, NomBank, and FrameNet annotations in order to recognize (and classify) the semantic dependencies of verbal predicates and predicate nominals.

Features derived from the *Textual Alignment* and *Partial Interpretation* modules are then sent to an *Entailment Computation* module, which uses a classifier (based on decision trees) in order to estimate the probability that a hypothesis is textually entailed by a text. For full details of the features used in this classifier, see (Hickl *et al.* 2006).

After classification, each entailment pair is sent to a final *Validation* module, which employs a set of syntactic heuristics (based on (Vanderwende *et al.* 2006)) in order to recognize instances of “false entailment” where a textual entail-



(a)



(b)

Figure 3: Hypothesis Generation.

ment relationship is known to not exist, despite either high confidence alignments and/or a positive entailment classification.

4. Generating Hypotheses from Text

Our system for RbE acquires knowledge from texts by recognizing entailment relationships between a series of text or knowledge probes and a series of hypotheses that have been automatically generated from a text corpus. In this section, we present a method for *hypothesis generation* which uses a Markov Chain model in order to generate natural language hypotheses which can be evaluated by an RTE or RKE system. Given the first text from Table 1, our method allows for the generation of the 10 hypotheses listed in Figure 4.

Text: Scientists have discovered that drinking tea protects against heart disease by improving the function of the artery walls.
Hyp₁: Tea protects from some diseases.
Hyp₂: Drinking tea is good for health.
Hyp₃: Scientists have linked the drinking of tea to improved health.
Hyp₄: Drinking tea can help maintain the elasticity of artery walls.
Hyp₅: Tea drinking studies were conducted.
Hyp₆: Drinking tea has been shown to prevent heart disease.
Hyp₇: Inelastic arteries are one symptom of heart disease.
Hyp₈: Heart disease is the leading cause of death in the U.S.
Hyp₉: 12.2 million Americans are diagnosed with heart disease.
Hyp₁₀: An estimated 2.3 billion people drink tea.

Figure 4: Examples of Generated Hypotheses.

Similar to the method for decomposing questions reported in (Harabagiu *et al.* 2006), we generate hypotheses using a sequence of operations $\{O_i\}$ which operate over sets of relations extracted from a text collection. Our current method uses two operations: (1) ADDITION, which appends a relation to a candidate hypothesis, and (2) REMOVAL, which subtracts a relation from a candidate hypothesis.

Figures 3(a) and 3(b) depict how these two relations can be used to generate *Hyp₁* and *Hyp₂* from Figure 4 respectively.

Our method for generating hypothesis operates on a Markov Chain (MC) by performing a random walk on a bipartite graph of (1) sequences of operators on relations between concepts related to the text, and (2) partial hypotheses created by the previous sequence of operators. The MC then alternates between selecting a sequence of operations $\{SO_i\}$ and generating a partial hypothesis H_i . The Markov Chain we use is illustrated in Figure 5.

We define the Markov Chain in the following way. First, we assume the initial state of the chain is equal to the initial sequence of operators $\{SO_0\}$ available to the Hypothesis Generation module. In order to select this sequence, we need to have access to a knowledge mapping function $M_1(KB, T, TC)$, where KB is equal to the available knowledge base, T is the text, and TC represents a set of text concepts, similar to the concepts depicted in Figures 3(a) and 3(b). We assume that SO_0 then represents a set of operators (where $|SO_0| < 4$) that maximizes the value of M_1 using dynamic programming. Here, we assume the role of M_1 is to coerce knowledge from a conceptual representation of a text, similar to the types of coercion functions used in metonymy resolution proposed by (Lapata and Lascarides 2003).

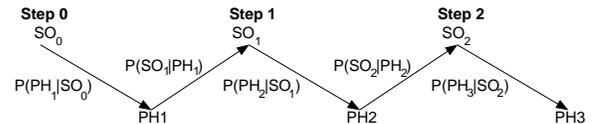


Figure 5: Markov Chain for Hypothesis Generation.

The state transition probabilities also depend on another knowledge mapping function M_1 which returns both (1) a list of related concepts from the KB ($\{C_L\}$) and (2) a list of relations between such concepts ($\{R_L\}$). For any given text T , we define M_1 as $M_1(KB, T) = \{C_L, R_L\}$.¹ We can now define a random walk for hypothesis generation using a matrix notation. Given $N = |C_L|$ and $M = |R_L|$, we consider A to be a $N \times M$ stochastic matrix with entries $a_{i,j} = p(r_i|h_j)$, where r_i is a relation sequence and h_j is a partial hypothesis. Similarly, we define another stochastic matrix

¹We anticipate that $\{C_L\}$ and $\{R_L\}$ are discovered by M_1 .

B (of dimensions $M \times N$) where $b_{i,j} = p(h_i|r_j)$. We estimate probabilities for $a_{i,j}$ and $b_{i,j}$ by applying the Viterbi algorithm to the maximum likelihood estimations resulting from the knowledge mappings for M_1 and M'_1 . We are currently studying several possibilities for these two functions, including approaches inspired by the density functions introduced by (Resnik 1995) and the more recent probabilistic framework for taxonomy representation described in (Snow et al. 2006).

5. Recognizing Knowledge Entailment

(Glickman & Dagan 2005) were inspired by probabilistic reasoning approaches when setting up their lexical textual entailment model. Our knowledge entailment model follows partially from this inspiration, but it also owes a great deal to work done in probabilistic abduction, as reported in (Raina et al. 2005). Abductive reasoning can be used to read from texts and to identify new knowledge that can be acquired and consolidated into a knowledge base.

In order to be able to infer a hypothesis from a text, we argue that knowledge entailment needs to rely on a model based on probabilistic abduction rather than one strictly based on first-order logic. In following the framework for weighted abduction first proposed in (Hobbs et al. 1993), (Raina et al. 2005) introduce a probabilistic abduction model which relaxes the unifications required by resolution refutation. This relaxation depends on a FOL representation which represents the knowledge derived from text in terms of lexicalized predicates and their arguments.² (Raina et al. 2005) only considered three possible relaxations in their work. Relaxations were considered (1) when the predicates considered were not identical, (2) when the considered predicates had different numbers of arguments, and (3) when some of the arguments considered were coreferential. Each of these three relaxations is then interpreted as an abductive assumption that has a plausibility score which is quantified using an *assumption cost model*.

In the method for knowledge entailment that we propose, the model for generating a hypothesis incorporates – via the random walk model for hypothesis generation proposed in the previous section – a probabilistic model of the plausibility of the hypothesis, given the concepts it uses and the relations that have been identified between them. If the TE system validates a positive entailment relationship between the text and the hypothesis, we expect that the probabilistic abduction employed by knowledge entailment will do so as well. This method allows us to focus on the knowledge support provided by the probabilistic abduction, instead of worrying about identifying the right abduction relaxations for our model, given the information derived from the text and hypothesis.

In our work, we have considered two forms of knowledge mappings for probabilistic abduction:

- $\mathbf{M}_1(\text{text}, \text{KB})$: A set of functions that estimates the cost of using knowledge *already encoded* in the knowledge base in order to perform abduction;

²In addition to (Raina et al. 2005)’s work, this representation was used in (Harabagiu et al. 2000) and (Moldovan et al. 2003).

- $\mathbf{M}_2(\text{text}, \text{KB})$: A set of functions that estimates the cost of discovering *new knowledge* in order to perform abduction

Following (Harabagiu et al. 2006), we anticipate that the cost functions associated with M_1 and M_2 can be derived from a mixture model that incorporates an estimate for the cost of abduction with the probability d of generating a particular hypothesis. We define the cost of abduction, $C(A_i)$, in the following way:

$$C(A_i) = d * C(M_1) + (1 - d)C(M_2) * C(A_{i-1})$$

The definition of these cost functions is an area of much current work for us. Although we have explored modeling the cost of performing abduction using existing knowledge (M_1) using world knowledge derived from lexical chains extracted from the sense-disambiguated glosses found in Extended WordNet³, we are actively investigating other forms of knowledge representation. In order to model the process of acquiring new knowledge to perform abduction (M_2), we have explored techniques for discovering new instances of categories based on work by (Downey et al. 2005) and (Etzioni et al. 2005).

No matter what mapping functions we identify for M_1 and M_2 , we assume that the cost of abduction will be directly dependent on the process of generating hypotheses for consideration in KE. Given k generated hypotheses, we expect that we can define a stochastic matrix A of dimensions $k \times k$ such that $a_{ij} = \alpha * Cost(A_i)$.

Given k hypotheses, we expect that we can define a stochastic matrix A of dimensions $k \times k$ such that $a_{ij} = \alpha * Cost(A_i, A_j)$. In a similar fashion, we expect that we can compute another stochastic matrix B of dimension $k \times k$ is defined such that $b_{ij} = \beta_j * M_2(i, j)$, where $\beta_j = 1 / \sum_{i=1}^k M_2(i, j)$.

With these two matrices in place, we then compute a relevance vector R for all generated hypotheses which can be defined by $R = [dA + (1 - d)B]_i \times R$. The resultant square matrix $E = [dA + (1 - d)B]$ then can be used to define a Markov Chain where each element e_{ij} from E specifies the transition probability from state i to state j in the chain. This allows us then to compute a value for d , i.e. the probability that a hypothesis is generated which can be evaluated in KE using knowledge available in the KB. In addition, this method also allows us to compute $(1 - d)$, the probability that the evaluation of KE depends on new knowledge which must be discovered and added to the KB.

We believe the unification of textual entailment with the abductive reasoning provides a framework that can be used to provide greater control over the knowledge being added to a KB. As depicted in Figure 6, the introduction of error-full knowledge can significantly degrade the quality of a KB as Machine Reading progresses. In addition to methods for detecting errors, systems need to include methods for evaluating the quality of previous knowledge base states. We anticipate that RKE systems will need to incorporate additional knowledge control functions – which we refer to as *Error Recovery (ER_i)* functions – which will allow for the identification of potential errors in the KB by estimating the cost

³Available at <http://xwn.hlt.utdallas.edu/>.

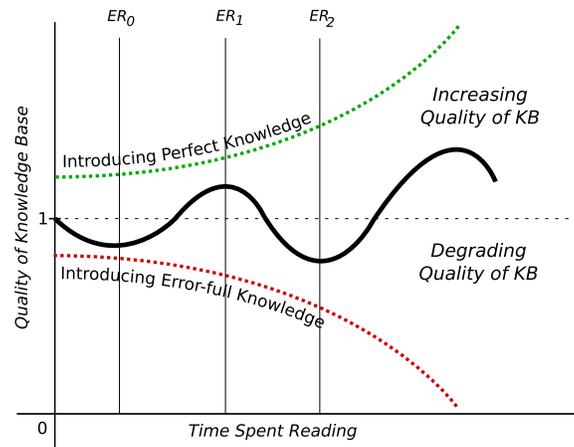


Figure 6: Ensuring Stability and Control.

of performing a particular abduction given any of the previous states of the KB. This would allow for the evolution of a KB over time to be cast as another stochastic process where the expected likelihood of adding new knowledge to the KB could be computed both in terms of the cost of performing the abduction as well as the probability that performing the abduction would lead to the introduction of errors.

6. Conclusions

This paper introduced Reading by Entailment (RbE), a new framework for both acquiring and consolidating knowledge derivable from text. We show that by leveraging two complementary systems for recognizing entailment relationships, we can reliably identify the set of hypotheses that can be inferred from a text corpus and ensure that additions to a KB remain consistent and controlled.

References

- Bar-Haim, R.; Dagan, I.; Dolan, B.; Ferro, L.; Giampiccolo, D.; Magnini, B.; and Szpektor, I. 2006. The Second PASCAL Recognising Textual Entailment Challenge. In *Proceedings of the Second PASCAL Challenges Workshop*.
- Bos, J.; and Makert, K. 2006. When logical inference determining textual entailment (and when it doesn't). In *Proceedings of the Second PASCAL Challenges Workshop*.
- Dagan, I.; Glickman, O.; and Magnini, B. 2006. The PASCAL Recognising Textual Entailment Challenge. In *MLCW 2005, LNAI Volume 3944*, pages 177-190.
- DelMonte, R.; Bristot, A.; Boniforti, M.; and Tonelli, S. 2006. Coping with semantic uncertainty with VENSES. In *Proceedings of the Second PASCAL Challenges Workshop*.
- de Marneffe, M.; MacCartney, B.; Grenager, T.; Cer, D.; Rafferty, A.; and Manning, C. 2006. Learning to distinguish valid textual entailments. In *Proceedings of the Second PASCAL Challenges Workshop*.
- Downey, D.; Etzioni, O.; and Soderland, S. 2005. A probabilistic model of redundancy for information extraction. In *Proceedings of IJCAI-2005*.
- Etzioni, O.; Cafarella, M.; Downey, D.; Kok, S.; Popescu, A.; Shaked, T.; Soderland, S.; Weld, D.; and Yates, A. 2005. Unsupervised named-entity extraction from the web: An experimental study. In *AI Journal*.
- Glickman, O.; and Dagan, I. 2005. A Probabilistic Setting and Lexical Co-occurrence Model for Textual Entailment. In *ACL Workshop on Empirical Modeling of Semantic Equivalence and Entailment*.
- Harabagiu, S.; Hickl, A.; and Lacatusu, F. 2006. Negation, Contrast, and Contradiction in Text Processing. In *AAAI-06*.
- Harabagiu, S.; Lacatusu, F.; and Hickl, A.; 2006. Answering Complex Questions with Random Walk Models. In *SIGIR-06*.
- Haghighi, A.; Ng, A.; and Manning, C. D. 2005. Robust textual inference via graph matching. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP-05)*.
- Hickl, A.; Williams, J.; Bensley, J.; Roberts, K.; Rink, B.; and Shi, Y. 2006. Recognizing Textual Entailment with LCC's Groundhog System. In *Proceedings of the Second PASCAL Challenges Workshop*.
- Moldovan, D.; Clark, C.; Harabagiu, S.; and Maiorano, S. 2003. COGEX: A Logic Prover for Question Answering. In *Proceedings of HLT/NAACL-2003*.
- Raina, R.; Ng, A.; and Manning, C. 2005. Robust textual inference via learning and abductive reasoning. In *Proceedings of AAAI 2005*.
- Tatu, M.; Iles, B.; Slavick, J.; Novischi, A.; and Moldovan, D. 2006. COGEX at the Second Recognizing Textual Entailment Challenge. In *Proceedings of the Second PASCAL Challenges Workshop*.
- Vanderwende, L.; Menezes, A.; and Snow, R. 2006. Microsoft Research at RTE-2: Syntactic Contributions in the Entailment Task: An Implementation. In *Proceedings of the Second PASCAL Challenges Workshop*.