Probabilistic Task Content Modeling for Episodic Textual Narratives

Eni Mustafaraj[1], **Martin Hoof**[2], **and Bernd Freisleben**[1]

[1]Dept. of Math. and Computer Science, University of Marburg, D-35032 Marburg, Germany
eni,freisleb@informatik.uni-marburg.de

[2]Dept. of Electrical Engineering, FH Kaiserslautern, D-67657 Kaiserslautern, Germany
m.hoof@et.fh-kl.de

Abstract

Episodic knowledge is often stored in the form of textual narratives written in natural language. However, a large repository of such narratives will contain both repetitive and novel knowledge. In this paper, we propose an approach for discovering interesting pieces of knowledge by using *a priori* task knowledge. By considering the narratives as generated by an underlying task structure, the elements of the task can be regarded as topics that generate the text. Then, by capturing task content in a probabilistic model, the model can be used, e.g., to identify the semantic orientation of textual phrases. An evaluation for a real world corpus of episodic narratives provides strong evidence for the feasibility of the proposed approach.

Episodic Textual Narratives

Discussing the use of knowledge in knowledge systems, (Richter 1998) distinguishes among three types of knowledge: background, contextual, and episodic knowledge. Out of all types of knowledge systems, episodic knowledge—which is of narrative character, because it tells the story of something that happened in the past—is directly employed in Case-Based Reasoning (CBR) systems only. Since episodic knowledge has a narrative character, a natural means of preserving it is in the form of textual narratives written in natural language by human users. Extraction of valuable pieces of knowledge from such narratives, which can serve as cases in the context of a Textual CBR (TCBR) system, is the focus of our research.

A common way of extracting knowledge from text documents in the context of a TCBR system is by considering *a priori* domain knowledge (Lenz 1999). The underlying idea is that text can be regarded as a container of domain objects (or information entities) and by using different types of knowledge acquisition, text documents can be reduced to a set of such information entities. In contrast, we take an alternative perspective. We consider text as generated by an underlying process, which consists of a series of related events. Domain objects are then participants of such events. By recognizing events and their participants in text, it is not only possible to discover domain objects, but also their respective roles in the event. In this way, text will not be a mere set of information entities, but a network of interconnected, semantically labeled entities. Additionally, a second difference to (Lenz 1999)'s approach is that instead of *a priori* domain knowledge, we use *a priori* task knowledge to perform knowledge extraction.

Task Knowledge: An Example

Consider the task MONITOR-and-DIAGNOSE, a task that starts with the monitoring of an object and continues with diagnosis only if something problematic is observed during the monitoring step. In general, there is a distinct temporal order in the way the MONITOR-and-DIAGNOSE task is performed.

1. Some entities of interest are observed and the respective findings are noticed.
2. These findings are then explained and evaluated.
3. If findings are evaluated as negative, actions for maintenance are recommended.

These three steps can be referred to as observation (`OBS`), explanation (`EXP`), and take action (`ACT`), and can be regarded as events that occur during the execution of the MONITOR-and-DIAGNOSE task. In the same way in which these events occur in reality, they will be described in written form, too. Each of the events is constituted by relations between different elements of the task. For example, an `OBS` relates an observed object to a finding, or an `EXP` relates a symptom to a possible cause. In our previous research, we have described an active learning approach (named LARC) that learns to annotate episodic narratives with task knowledge roles such as observed object, finding, cause, etc. (Mustafaraj, Hoof, & Freisleben 2006).

Probabilistic Task Content Modeling

The previous description of a task as a series of interconnected events and participant roles constitutes an abstract model of task structure. An instantiation of the task structure in a real situation produces the task content. The instantiation consists of the verbalization of the abstract events and roles with concrete sentences and phrases. Because real-world events and natural language are of stochastic nature, a probabilistic model is an appropriate means for capturing task content.

An approach for building and using probabilistic content modeling has been proposed in (Barzilay & Lee 2004). The assumption, upon which the authors build their approach, is that text documents from a specific domain display high similarity, since they are generated from the same content model. In representing a content model with a Hidden Markov Model, it is possible to interpret the states of the models as topics that generate the sentences of the text, and state transitions as topic shifts in the content. The authors have successfully used the learned models for the tasks of information ordering and extractive summarization. However, the learned topics do not have an explicit meaning, because the model is built by using unlabeled data.

In the context of episodic textual narratives of task content, where *a priori* task knowledge is available, it is possible to hypothesize a probabilistic model structure based on the elements of task structure and use available data to learn the parameters of the model. The fact that the corpus of narratives has been previously annotated with event types and roles by LARC permits the use of the model for different types of inquiries. One type that we investigate in this paper is the determination of semantic orientation for textual phrases annotated with the role of `finding` in the MONITOR-and-DIAGNOSE task.

Semantic Orientation of Phrases

Semantic orientation is concerned with the polarity of phrases, that is, whether the meaning of a phrase is positive or negative. A related problem is sentiment analysis, concerned with the polarity of a whole article (e.g. book, movie, or other product reviews) (Cui, Mittal, & Datar 2006).

In the MONITOR-and-DIAGNOSE task, an important piece of knowledge is that annotated with the `finding` role, because it will be used for indexing. Based on the condition of an observed object, a finding can be positive or negative. Some examples of phrases with negative orientation are: "higher values", "clear deviations"; while phrases with nonnegative orientation are: "in normal area", "insignificant change", etc. Our goal is to automatically identify which are negative and positive finding phrases.

The knowledge captured in the probabilistic task content model (PTCM) (that is built by following only in part (Barzilay & Lee 2004), since the states are known in our case) can be used to create two language models:

$$P(+) = P(phrase|FI, \phi) \quad (1)$$

$$P(-) = P(phrase|FI, AC) \quad (2)$$

Equation 1 is the language model of the positive (+) phrases, which are emmited from the the PTCM model when the (hidden) state `FI` (finding) is followed by the empty state ϕ. Equation 2 is the language model for negative (-) phrases, which are emmited by the state `FI` followed by the state `AC` (action). This definition for the two models is based on task knowledge, because only negative findings will be followed by action reccommendations. All the other phrases that are not automatically assigned to one of the models, will be classified by the PTCM classifier that selects for each phrase the model with the higher score.

Evaluation

Our corpus of narratives consists of descriptions of performing the task MONITOR-and-DIAGNOSE in the field of diagnosis for electrical machines. More details on the corpus are found in (Mustafaraj, Hoof, & Freisleben 2006). We evaluate here the performance of the PTCM classifier in the task of identifying the semantic orientation of textual phrases. For comparison purposes, we built two others classifiers. A simple baseline classifier assigns labels randomly (by coin toss). Then, as an informed baseline, we have built a classifier based on domain-specific lexical knowledge. For that, all single words that contain polarity orientation (within the given context) were manually labeled as positive (e.g. 'normal', 'low'') or as negative (e.g. 'irregular', 'scattered'). The classifier performs a majority vote among the words of a phrase; in case of ties, it decides for the positive class (the most frequent). The average F-measures (the harmonic mean of the precision and recall metrics) for the three classifiers calculated over 10 trials are shown in Figure 1. The PTCM classifier performs significantly better than the informed classifier, according to the Wilcoxon signed-rank test ($p < 0.05$).

Classifer	F-measure
Random Classifier	0.47
Informed Classifier	0.76
PTCM Classifier	0.83

Table 1: F-measure for three different classifiers.

Conclusions

In this paper, we provided empirical evidence on how a probabilistic model, which was built by using task knowledge only, was able to identify the semantic orientation of textual phrases at a significantly higher level that an informed classifier (containing domain-specific lexical knowledge). Identifying the semantic orientation of textual phrases is important in the context of a TCBR system, in order to distinguish between interesting and uninteresting cases.

References

Barzilay, R., and Lee, L. 2004. Catching the drift: Probabilistic content models, with applications to generation and summarization. In *Proc. of NAACL-HLT*, 113–120.

Cui, H.; Mittal, V.; and Datar, M. 2006. Comparative experiments on sentiment classification for online product reviews. In *Proc. of AAAI*, 1265–1270.

Lenz, M. 1999. *Case Retrieval Nets as a Model for Building Flexible Information Systems*. Humboldt Universitaet zu Berlin.

Mustafaraj, E.; Hoof, M.; and Freisleben, B. 2006. LARC: Learning to assign knowledge roles to textual cases. In *Proc. of FLAIRS*, 370–375.

Richter, M. 1998. Introduction. In *Case-Based Reasoning Technology: From Foundations to Applications*, 1–15. Springer Verlag.