Automatic Detection of Causal Relations for Question Answering

Roxana Girju
Computer Science Department
Baylor University
Waco, Texas
roxana@cs.baylor.edu

Abstract

Causation relations are a pervasive feature of human language. Despite this, the automatic acquisition of causal information in text has proved to be a difficult task in NLP. This paper provides a method for the automatic detection and extraction of causal relations. We also present an inductive learning approach to the automatic discovery of lexical and semantic constraints necessary in the disambiguation of causal relations that are then used in question answering. We devised a classification of causal questions and tested the procedure on a QA system.

1 Introduction

The automatic detection of semantic information in English texts is a very difficult task, as English is highly ambiguous. However, there are many applications which can greatly benefit from in depth semantic analysis of text. Question Answering is one of them.

An important semantic relation for many applications is the *causation relation*. Although many computational linguists focused their attention on this semantic relation, they used hand-coded patterns to extract causation information from text.

This work has been motivated by our desire to analyze cause-effect questions that are currently beyond the state-of-the-art in QA technology. This paper provides an inductive learning approach to the automatic discovery of lexical and semantic constraints necessary in the disambiguation of verbal causal relations. After a brief review of the previous work in Computational Linguistics on causation in section 2, we present in section 3 a classification of lexico-syntactic patterns that are used to express causation in English texts and show the difficulties involved in the automatic detection and extraction of these patterns. A method for automatic detection of causation patterns and validation of ambiguous verbal lexico-syntactic patterns referring to causation is proposed in section 4. Results are discussed in section 5, and in section 6 the application of causal relations in Question Answering is demonstrated.

2 Previous Work in Computational Linguistics

Computational linguists have tried to tackle the notion of causality in natural language focusing on lexical and semantic constructions that can express this relation.

Many previous studies have attempted to extract implicit inter-sentential cause-effect relations from text using knowledge-based inferences (Joskowiscz et al. 1989), (Kaplan 1991). These studies were based on hand-coded, domain-specific knowledge bases difficult to scale up for realistic applications. More recently, other researchers (Garcia 1997), (Khoo et al. 2000) used linguistic patterns to identify explicit causation relations in text without any knowledge-based inference. Garcia used French texts to capture causation relationships through linguistic indicators organized in a semantic model which classifies causative verbal patterns. She found

25 causal relations with an approach based on the "Force Dynamics" of Leonard Talmy claiming a precision of 85%.

Khoo at al. used predefined verbal linguistic patterns to extract cause-effect information from business and medical newspaper texts. They presented a simple computational method based on a set of partially parsed linguistic patterns that usually indicate the presence of a causal relationship. The relationships were determined by exact matching on text with a precision of about 68%.

3 How are causation relations expressed in English?

Any causative construction involves two components, the cause and its effect. For example:
"The bus fails to turn up. `As a result`, I am late for a meeting".(Comrie 1981)

Here the cause is represented by the bus's failing to turn up, and the effect by my being late for the meeting.

In English, the causative constructions can be **explicit** or **implicit**. Usually, explicit causation patterns can contain relevant keywords such as *cause, effect, consequence*, but also ambiguous ones such as *generate, induce*, etc. The implicit causative constructions are more complex, involving inference based on semantic analysis and background knowledge. The English language provides a multitude of cause-effect expressions that are very productive. In this paper we focus on explicit but ambiguous verbal causation patterns and provide a detailed computational analysis. A list of other causation expressions were presented in detail elsewhere (Girju 2002).

Causation verbs

Many linguists focused their attention on causative verbal constructions that can be classified based on a lexical decomposition. This decomposition builds a taxonomy of causative verbs according to whether they define only the causal link or the causal link plus other components of the two entities that are causally related (Nedjalkov and Silnickij 1969):

1. **Simple causatives** (*cause*, *lead to*, *bring about*, *generate*, *make*, *force*, *allow*, etc.)
Here the linking verb refers only to the causal link, being synonymous with the verb *cause*. E.g., "Earthquakes `generate` tidal waves."

2. **Resultative causatives** (*kill*, *melt*, *dry*, etc.)
These verbs refer to the causal link plus a part of the resulting situation.

3. **Instrumental causatives** (*poison* (killing by poisoning), *hang*, *punch*, *clean*, etc.)
These causatives express a part of the causing event as well as the result.

4 Automatic detection of causation relationships

In this section we describe a method for automatic detection of lexico-syntactic patterns that express causation.

The algorithm consists of two major procedures. The first procedure discovers lexico-syntactic patterns that can express the causation relation, and the second procedure presents an inductive learning approach to the automatic detection of syntactic and semantic constraints on the constituent components.

4.1 Automatic discovery of lexico-syntactic patterns referring to causation

One of the most frequent explicit intra-sentential patterns that can express causation is $<NP_1\ verb\ NP_2>$. In this paper we focus on this kind of patterns, where the verb is a simple causative.

In order to catch the most frequently used lexico-syntactic patterns referring to causation, we used the following procedure (Hearst 1998):

Discovery of lexico-syntactic patterns:
`Input`: semantic relation R
`Output`: lexico-syntactic patterns expressing R

STEP 1. Pick a semantic relation R (in this paper, CAUSATION)
STEP 2. Pick a pair of noun phrases C_i, C_j among which R holds.

Since CAUSE-TO is one of the semantic relations explicitly used in WordNet, this is an excellent resource for picking C_i and C_j. The CAUSE-TO relation is a transitive relation between verb synsets. For example, in WordNet the second sense of the verb *develop* is "*causes to grow*". Although WordNet contains numerous causation relationships between nouns that are always true, they are not directly mentioned. One way to determine such relationships is to look for all patterns $<NP_1\ cause\_vb\ NP_2>$ that occur between a noun entry and another noun in the corresponding gloss definition. One such example is the causation relationship between $\{$*bonyness*$\}$ and $\{$*starvation*$\}$. The gloss of $\{$*bonyness* (#1/1)$\}$ is (*extreme leanness (usually caused by starvation or disease)*).

WordNet 1.7 contains 429 such relations linking nouns from different domains, the most frequent being medicine (about 58.28%).

STEP 3. Extract lexico-syntactic patterns that link the two selected noun phrases by searching a collection of texts.

For each pair of causation nouns determined above, search the Internet or any other collection of documents and retain only the sentences containing the pair. From these sentences, determine automatically all the patterns $<NP_1$ *verb/verb_expression* $NP_2>$, where $<NP_1$ - $NP_2>$ is the pair considered.

The result is a list of verbs/verbal expressions that refer to causation (see Table 1). Some of these verbs are always referring to causation, but most of them are ambiguous, as they express a causation relation only in a particular context and only between specific pairs of nouns. For example, $<NP_1$ *produces* $NP_2>$. In most cases, the verb *produce* has the sense of *manufacture*, and only in some particular contexts it refers to causation.

In this approach, the acquisition of linguistic patterns is done automatically, as the pattern is predefined ($<NP_1$ *verb* $NP_2>$). As described in the next subsections, the relationships are disambiguated and only those referring to causation are retained.

4.2 Learning Syntactic and Semantic Constraints for causal relation

The learning procedure proposed here is *supervised*, for the learning algorithm is provided with a set of inputs along with the corresponding set of correct outputs. Based on a set of positive and negative causal training examples provided and annotated by the user, the algorithm creates a decision tree and a set of rules that classify new data. The rules produce constraints on the noun constituents of the lexical patterns.

For the discovery of the semantic constraints we used C4.5 *decision tree learning* (Quinlan 1999). The learned function is represented by a decision tree, or a set of if-then rules. The decision tree learning searches a complete hypothesis space from simple to complex hypotheses until it finds a hypothesis consistent with the data. Its bias is a preference for the shorter tree that places high information gain attributes closer to the root.

The error in the training examples can be overcome by using different training and a test corpora, or by cross-validation techniques.

C4.5 receives in general two input files, the NAMES file defining the names of the attributes, attribute values and classes, and the DATA file containing the examples.

4.2.1 Preprocessing Causal Lexico-Syntactic Patterns

Since a part of our constraint learning procedure is based on the semantic information provided by WordNet, we need to preprocess the noun phrases (NPs) extracted and identify the *cause* and the *effect*. For each NP we keep only the largest word sequence (from left to right) that is defined in WordNet as a concept.

For example, from the noun phrase "*a 7.1 magnitude earthquake*" the procedure retains only "*earthquake*", as it is the WordNet concept with the largest number of words in the noun phrase.

We did not consider those noun phrases in which the head word had other part of speech than noun.

4.2.2 Building the Training Corpus and the Test Corpus

In order to learn the constraints, we used the LA TIMES section of the TREC 9 text collection. For each of the 60 verbs generated with the procedure described in section 4.1, we searched the text collection and retained 120 sentences containing the verb. Thus, a training corpus "A" of 6,000 sentences, and respectively, a test corpus of 1,200 sentences were automatically created. Each sentence in these corpora was then parsed using the syntactic parser developed by Charniak (Charniak 1999).

Focusing only on the sentences containing relations indicated by the pattern considered, we manually annotated all instances matched by the pattern as referring to causation or not. Using the training corpus, the system extracted 6,523 relationships of the type $<NP_1$ *verb* $NP_2>$, from which 2,101 were

Causal verbs			
give rise (to)	stir up	create	start
induce	entail	launch	make
produce	contribute (to)	develop	begin
generate	set up	bring	rise
effect	trigger off	stimulate	
bring about	commence	call forth	
provoke	set off	unleash	
arouse	set in motion	effectuate	
elicit	bring on	kick up	
lead (to)	conduce (to)	give birth (to)	
trigger	educe		
derive (from)	originate in	call down	
associate (with)	lead off	put forward	
relate (to)	spark	cause	
link (to)	spark off		
stem (from)	evoke		
originate	link up		
bring forth	implicate (in)		
lead up	activate		
trigger off	actuate		
bring on	kindle		
result (from)	fire up		

Table 1: Ambiguous causation verbs detected with the procedure described in section 4.1.

causal relations, while 4,422 were not.

4.2.3 Selecting features

The next step consists of detecting the constraints necessary on nouns and verb for the pattern $<NP_1$ *verb* $NP_2>$ such that the lexico-syntactic pattern indicates a causation relationship.

The basic idea we employ here is that only some categories of noun phrases can be associated with a causation link. According to the philosophy researcher Jaegwon Kim (Kim 1993), any discussion of causation implies an ontological framework of entities among which causal relations are to hold, and also "an accompanying logical and semantical framework in which these entities can be talked about". He argues that the entities that represent either causes or effects are often *events*, but also *conditions, states, phenomena, processes*, and sometimes even *facts*, and that coherent causal talk is possible only within a coherent ontological framework of such states of affairs.

Many researchers ((Blaheta and Charniak 2000), (Gildea and Jurafsky 2000), showed that lexical and syntactic information is very useful for predicate-argument recognition tasks, such as semantic roles. However, lexical and syntactic information alone is not sufficient for the detection of complex semantic relations, such as CAUSE.

Based on these considerents and on our observations of the English texts, we selected a list of 19 features which are divided here into two categories: *lexical* and *semantic* features.

The lexical feature is represented by the causation verb in the pattern considered. As verb senses in WordNet are fine grained providing a large list of semantic hierarchies the verb can belong to, we decided to use only the lexical information the verb provides. The values of this feature are represented by the 60 verbs detected with the procedure described in section 4.1. This feature is very important, as our intention here is to capture the semantic information brought by the verb in combination with the subject and object noun phrases that attach to it.

As we don't use word sense disambiguation to disambiguate each noun phrase in context, we have to take into consideration all the WordNet semantic hierarchies they belong to according to each sense. For each noun phrase representing the cause, and respectively the effect, we used as semantic features the 9 noun hierarchies in WordNet: *entity, psychological feature, abstraction, state, event, act, group, possession,* and *phenomenon*. Each feature is true if it is one of the semantic classes the noun phrase can belong to, and false otherwise.

4.2.4 Learning Algorithm

`Input:` positive and negative causal examples

`Output:` lexical and semantic constraints

Step 1. Generalize the training examples
Initially, the training corpus consists of examples that contain only lexical features in the following format:

```
<cause_NP; verb; effect_NP;
         target>,
```

where *target* can be either "Yes" or "No", depending whether or not an example encodes cause.

For example, `<earthquake; generate; Tsunami; Yes>` indicates that between the noun "*earthquake*" and the noun "*Tsunami*" there is a cause relation.

From this intermediate corpus a generalized set of training examples was built, by expanding each intermediate example with the list of semantic features using the following format:

```
<entityNP1,
psychological-featureNP1,
abstractionNP1, stateNP1,
eventNP1, actNP1, groupNP1,
possessionNP1, phenomenonNP1;
verb;
entityNP2,
psychological-featureNP2,
abstractionNP2, stateNP2,
eventNP2, actNP2, groupNP2,
possessionNP2, phenomenonNP2;
target>.
```

For instance, the initial example becomes `<f, f, f, f, f, f, f, f, t, generate, f, f, f, f, f, t, f, f, f, yes>`, as the noun phrase *earthquake* belongs only to the `{phenomenon}` noun hierarchy and the noun phrase *Tsunami* is only in the `{event}` noun hierarchy in WordNet.

Step 2. Learning constraints from training examples
For the examples in the generalized training corpus (those that are either positive or negative), constraints are determined using C4.5. In this context, the features are the characteristics that distinguish the causal relation, and the values of the features are either specific words (e.g., the verb) or their WordNet corresponding semantic classes (the furthest ancestors in WordNet of the corresponding concept).

On this training corpus we applied C4.5 using a 10-fold cross validation. The output is represented by 10 sets of rules generated from the positive and negative examples.

The rules in each set were ranked according to their frequency of occurrence and average accuracy obtained for that particular set. In order to use the best rules, we decided to keep only the ones that had a frequency above a threshold (occur in at least 7 of the 10 sets of rules) and with an average accuracy greater than 60%.

4.2.5 The Constraints

Table 2 summarizes the constraints learned by the program.

As we can notice, the constraints combine information about the semantic classes of the noun phrases representing the cause and effect with the lexical information about the verb.

5 Results

To validate the constraints for extracting causal relations, we used the test corpus "B".

For each head of the noun phrases in the CAUSE and EFFECT positions, the system determined automatically the most general subsumers in WordNet for each sense. The test corpus contained 683 relationships of the type $<NP_1$ *verb* $NP_2>$, from which only 115 were causal patterns. The results provided by the causal relation discovery procedure were validated by a human annotator.

Let us define the *precision* and *recall* performance metrics in this context.

$$precision = \frac{Number\ of\ correctly\ retrieved\ relations}{Number\ of\ relations\ retrieved}$$

$$recall = \frac{Number\ of\ correctly\ retrieved\ relations}{Number\ of\ correct\ relations}$$

The system retrieved 138 relations, of which 102 were causal relations and 36 were non-causal relations, yielding a precision of 73.91% and a recall of 88.69%. Table 3 shows the results obtained for the pattern considered.

However, there were other 38 causal relations found in the corpus, expressed by other than the lexico-syntactic pattern considered in this paper,

Nr	Class-NP1	verb	Class-NP2	Target	Acc.(%)	Freq.	Example
0	*	cause	*	1	100	18	*hunger causes headache*
1	*	*	phenomenon	1	98	38	*movement triggers earthquake*
2	!entity	associated-with or related-to	!abstraction and !group and !possession	1	63.00	26	*syndromes are associated with disease*
3	!entity	*	event	1	89	24	*inactivation induces events*
4	!abstraction	*	event or act	1	90	12	*event generated group action*
5	*	lead-to	!entity and !group	1	88	21	*intake leads to immunodeficiency*
6	*	induce	entity or abstraction	0	70.0	10	*carcinogens induce fields*
7	*	*	!state and !event and !act and group	0	70.7	10	*path leads to house*
8	entity	*	!state and !event and !phenomenon	0	70.0	10	*cells derived from lymph nodes*

Table 2: The list of constrains accompanied by examples (! means "is not", 1 means "Is a causal relation,", 0 means "Is not a causal relation", and * means anything)

yielding a global causal relation coverage (recall) of 66.6% [102/115+38].

The errors are explained mostly by the fact that the causal pattern is very ambiguous. This lexico-syntactic pattern encode numerous relations which are very difficult to disambiguate based only on the list of connectors.

The errors were also caused by the incorrect parsing of noun phrases, the use of the rules with smaller accuracy (e.g. 63%), and the lack of named entities recognition in WordNet (e.g., names of people, places, etc.).

Some of the factors that contributed a lot to the precision and recall results were the size and the accuracy of the positive and negative examples in the training corpus. For this experiment we used only a fairly small training corpus of 6,523 examples.

No. of Relations	**Causal pattern** $<NP_1\ verb\ NP_2>$
Number of patterns	683
Number of correct relations	115
Number of relations retrieved	138
Number of correctly retrieved relations	102
Precision	73.91%
Recall	88.69%

Table 3: The number of relations obtained and the accuracy for the causal pattern used for this research.

6 Importance and application of causal relations in Question Answering

Causation relationships are very pervasive, but most of the time they are ambiguous or implicit. The degree of ambiguity of these relations varies with the semantic possibilities of interpretation of the constituent syntactic terms. This disambiguation proves to be very useful for applications like Question Answering.

Causation questions can be mainly introduced by the following question types: *what, which, name* (what causes/be the cause of, what be the effect of, what happens when/after, what cause_vb <object>), *how* (how <causation_adj>), and *why*. However, an analysis of these question types alone is not sufficient for causation, another classification criteria being required. Based on our observation of cause-effect questions, we propose the following question classes based on their ambiguity:

1. **Explicit causation questions**
The question contains explicit unambiguous keywords that define the type of relation, and determines the semantic type of the question (e.g., *effect, cause, consequence*, etc.)

"What are the `causes` of lung cancer?"

"Name the `effects` of radiation on health."

"Which were the `consequences` of Mt. Saint Elena eruption on fish?"

2. **Ambiguous (semi-explicit) causation questions**
The question contains explicit but ambiguous keywords that refer to the causation relation. Once disambiguated, they help in the detection of the semantic type of the question (e.g., *lead to, produce, generate, trigger, create*, etc.)

"Does watching violent cartoons `create` aggression in children?"

"What economic events `led to` the extreme wealth among Americans in the early 1920's?"

3. **Implicit causation questions**
This type of questions involves reasoning, based on deep semantic analysis and background knowledge. They are usually introduced by the semantic types *why*, *what*, *how*, and can be further classified in two important subtypes:

a) *Causation questions disambiguated based on the semantic analysis of question keywords*
"`Why` did Socrates `die`?"
"`What killed` Socrates?"
"How `dangerous` is a volcanic eruption?"
"Is exercise `good` to the brain?"

It is recognized that questions of type *what*, and even *how* and *why*, are ambiguous, and usually the question is disambiguated by other keywords in the question.

In the example question "What killed Socrates?", the verb *kill* is a causation verb meaning *cause to die*, so the second question asks for the cause of the Socrates' death.

The *why* questions are more complex asking for explanations or justifications. Explanations can be expressed in English in different ways, not always referring to causation. Thus, it is very difficult to determine directly from the question what kind of information we should look for in the answer.

b) *Causation questions that are disambiguated based on how the answer is expressed in the text*

Behavioral psychologists illustrated that there are several different ways of answering *why* questions in biology. For example, the question "*Why do robins sing in the spring?*" can have multiple categories of answers:
Causation. (What is the cause?)
Answer: "Robins sing in spring because increases in day length trigger hormonal action".
Development. (How does it develop?)
Answer: "Robins sing in spring because they have learned songs from their fathers and neighbors."
Origin. (How did it evolve?)
Answer: "Song evolved as a means of communication early in the avian lineage".
Function. (What is the function?)
Answer: "Robins sing in spring to attract mates."

The algorithm for automatic extraction of causation relations presented in section 4 was tested on a list of 50 natural language causation questions (50 explicit and 50 ambiguous) using a state-of-the-art Question Answering system (Harabagiu et al. 2001). The questions were representative for the first two categories of causation questions presented above, namely *explicit* and *ambiguous causation questions*. We selected for this purpose the TREC9 text collection and we (semi-automatically) searched it for 50 distinct relationships of the type $<NP_1\ verb\ NP_2>$, where the verb was one of the 60 causal verbs considered. For each such relationship we formulated a cause-effect question of the first two types presented above. We also made sure each question had the answer in the documents generated by the IR module.

Table 4 shows two examples of questions from each class. We also considered as good answer any other correct answer different from the one represented by the causal pattern. However these answers were not taken into consideration in the precision calculation of the QA system with the causation module included. The rational was that we wanted to measure only the contribution of the causal relations method. The 50 questions were tested on the QA system with (61% precision) and without (36% precision) the causation module included, with a gain in precision of 25%.

7 Discussion and Conclusions

The approach presented in this paper for the detection and validation of causation patterns is a novel one. Other authors (Khoo et al. 2000) restricted their text corpus to a medical/business database and used hand-coded causation patterns that were mostly unambiguous. Our method discovers automatically generally applicable lexico-syntactic patterns referring to causation and disambiguates the causation relationships obtained from the pattern application on

Question Class	Question	Answer	
		QA without causation module	**QA with causation module**
Explicit	What `causes` post-traumatic stress disorder?	Post-traumatic Stress Disorder - What are the Symptoms and Causes?	Post-traumatic stress disorder `results from` a traumatic event.
	What are the `effects` of acid rain?	Projects, reports, and information about the effects of acid rain	Acid rain is known to `contribute to` the corrosion of metals..
Ambiguous	What can `trigger` an allergic reaction?	The protein is consistent with something that triggers an allergic reaction	An antigen `producing` an allergic reaction is defined as an allergen.
	What phenomenon is `with` volcanoes?	.. that deglaciation are *associated* associated with increased volcanic activity..	There are often earthquakes `generated by` volcanism..

Table 4: Examples of cause-effect questions tested on a Question Answering system.

text. Moreover, we showed that the automatic detection of causal relations is very important in Question Answering for answering cause-effect questions.

References

D. Blaheta and E. Charniak, Assigning Function Tags to Parsed Text. In *Proceedings of the 1st Annual Meeting of the North American Chapter of the Association for Computational Linguistics*, Seattle, May 2000, pp. 234–240.

E. Charniak, A maximum-entropy-inspired parser. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL 2000)*, Seattle, WA.

B. Comrie. Causative constructions In *Language Universals and Linguistic Typology*, University of Chicago Press, Chicago, 1981.

S. Harabagiu, D. Moldovan, M. Pasca, M. Surdeanu, R. Mihalcea, R. Girju, V. Rus, F. Lacatusu, P. Moraescu, and R. Bunescu. 2001. Answering Complex, List and Context Questions with LCC's Question-Answering Server. In Proceedings of the *TExt Retrieval Conference for Question Answering (TREC 10)*.

M. Hearst. Automated Discovery of WordNet Relations. In *WordNet: An Electronic Lexical Database and Some of its Applications*, editor Fellbaum, C., MIT Press, 1998.

D. Garcia. COATIS, an NLP system to locate expressions of actions connected by causality links. In *Knowledge Acquisition, Modeling and Mangement*, The Tenth European Workshop, 1997.

D. Gildea and D. Jurafsky. Automatic Labeling of Semantic Roles. In *Proceedings of the 38th Annual Conference of the Association for Computational Linguistics (ACL-00)*, pages 512-520, Hong Kong, October 2000.

R. Girju. Text Mining for Semantic Relations. *Ph.D. Dissertation*, University of Texas at Dallas, May 2002.

L. Joskowiscz, T. Ksiezyk and R. Grishman. Deep domain models for discourse anaysis. In *The Annual AI Systems in Government Conference*.

R.M. Kaplan, and G. Berry-Rogghe. Knowledge-based acquisition of causal relationships in text. In *Knowledge Acquisition*, 3(3), 1991.

C. Khoo, S. Chan and Y. Niu. Extracting Causal Knowledge from a Medical Database Using Graphical Patterns In *Proceedings of ACL*, Hong Kong, 2000.

J. Kim. Causes and Events: Mackie on Causation. In *Causation*, ed. Ernest Sosa, and Michael Tooley, Oxford University Press, 1993.

V.P. Nedjalkov and G. Silnickij. The topology of causative constructions. In *Folia Linguistica* (6).

J.R. Quinlan. C4.5: Programs for Machine Learning. *Morgan Kaufmann*.