Mining Answers for Causation Questions

Roxana Girju and Dan Moldovan
University of Texas at Dallas
Dallas, Texas
roxana@utdallas.edu, *moldovan@utdallas.edu*

Abstract

Causal relations of various kinds are a pervasive feature of human language and theorising about the world. Despite this, the specification of a satisfactory general analysis of causal relations has long proved difficult. This paper provides a syntactic and semantic classification of cause-effect lexico-syntactic patterns for automatic detection and extraction of causation relationships in English texts. We also present a semi-automatic method of discovering generally applicable lexico-syntactic patterns that refer to the causation relation. The patterns are found automatically, but their validation is done semi-automatically.
Our final purpose is to add a new module to our existing Question Answering (QA) system that will answer complex cause-effect questions.

1. Introduction

The automatic extraction of linguistic patterns used to extract information relevant to a particular task from a collection of documents has captured the attention of many natural language processing researchers in the last decade. The automatic identification of semantic relations in text has become increasingly important in Information Extraction, Question Answering and Information Retrieval.

At the MUC competitions many Information Extraction systems used new and innovative techniques to discover relevant information from texts. In order to extract the exact answer to user queries, QA systems often need to synthesize information gathered from multiple documents or to identify new relationships between facts/entities and discover new knowledge.

An important semantic relation for all these 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 paper provides an analysis of causative construction representations in English texts, involving the interaction of various linguistic components, including semantics, syntax and morphology. In the next chapter we present different approaches of causation with emphasis on Artificial Intelligence and Computational Linguistics. Section 3 gives a detailed classification of lexico-syntactic patterns that are used to express causation in English texts. In section 4 we show the difficulties involved in the automatic detection and extraction of causation relations in text and we propose a method for automatic detection of causation patterns and a semi-automatic validation of ambiguous verbal lexico-syntactic patterns referring to causation. Results are discussed in section 5 and in section 6 we present a classification of causation questions and prove the importance of our approach on a Question Answering application. At the end we offer some discussion and conclusions.

2. Previous Work in Artificial Intelligence and Computational Linguistics

Broadly speaking, causality refers to the way of knowing if one state of affairs causes another. Although the notion of causality is very old (beginning with the Aristotle's *Metaphysics*), over the time it has been surrounded by controversy as scientists and philosophers have not agreed on the definition of causality and when two states of affairs are causally linked.

The theory of causality is very broad, and perhaps the most interesting feature of the work on causation over the last decades has been its diversity. Several theories have been developed resulting in an overwhelming number of publications. This explosion of approaches can be explained in part by the plurality of perspectives the researchers used, and by the diversity of domains to which the causation notion applies: ***philosophy, statistics, linguistics, physics, economics, biology, medicine***, etc.

According to Sowa (Sowa 2000), Artificial Intelligence is one of the three academic disciplines, besides theoretical physics and philosophy, that have addressed multiple and interesting questions about causality, developing theories intended to stimulate intelligent behavior at the human level and beyond. Much of the research on causation in AI has been done in *planning, explanation*, and *linguistic analysis*.

Planning in AI is the problem of finding a sequence of primitive actions to achieve some goal. The ability to reason about time and actions is fundamental to almost any intelligent entity that needs to make a series of decisions. However, it is difficult to represent the concept of taking some actions and the concept of the consequences of taking a series of actions, without dealing with the notions of time and causality. Planning actions for robots requires reasoning about the causal order of actions and about how much time it will take to perform the actions. Determining the cause of a certain state of affairs implies considering temporal precedence.

Although discussed in the context of robotics, planning is also important in many other areas of AI. In Natural Language Understanding, for instance, it is important to reason about peoples' plans and goals in order to best make sense of what they say.

Researchers in Natural Language Understanding ((Wilenski 1978), (Lim 1992)) have become increasingly aware of the similarities between the task of planning the content and structure of the natural language text and that of other AI planning tasks. Problems in NLP such as story, scene understanding, and language generation can be viewed as planning problems.

Explanation in AI deals with commonsense reasoning for rational actions, including causation, namely what defines the causal context and what differentiates it from other situations (Ortiz 1999), (Pearl 2000), (Hobbs 2001). The analysis of causation in this area is mostly done through the use of counterfactuals. Pearl defines his causal explanation as Bayesian networks, where links have an intrinsic directionality. His work focused on interpreting equations that express causal claims as claims about the outcomes of hypothetical experiments.

Ortiz examined in depth the role of the counterfactual reasoning in the theory of causation and the kind of inferencing that can be drawn in the course of causal attribution. He proposes a commonsense causal language and defines causation in terms of changes involved by the counterfactuals.

The discussion over the last twenty years brought some clouds on the adequacy of any singular causation analysis in terms of counterfactuals. However, lately there could be seen different refinements of how to achieve a closer match with commonsense reasoning about causation.

Computational Linguists have tried to treat the notion of causality in natural language focusing on lexical and semantic constructions that can express causation. Starting from the philosophical idea that causation relationships hold between two states of affairs, many linguists defined standard descriptions of causatives based on formal distinctions as to how the meaning of *causing* is indicated, and on a small set of descriptive labels indicating semantic constructions (Talmy 1976).

Many previous studies have attempted to extract implicit inter-sentential cause-effect relations from text using knowledge-based inferences (Joskowiscz, Ksiezyk and Grishman 1989), (Kaplan 1991). These studies were based on hand-coded, domain-specific knowledge bases difficult to scale up for realistic applications.

Other researchers (Garcia 1997), (Khoo et al. 2000) used linguistic patterns to identify explicitly expressed 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. The precision Garcia claims is 85%.

Khoo at al. used predefined verbal linguistic patterns to extract cause-effect information from business (1999) and medical newspaper texts (2001). They used a simple computational method with no knowledge-based inferencing and partial parsing of sentences, based on a set of 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?

Causative constructions have played an important role in the recent history of linguistics mainly because their study involves the interaction of various components of over-all linguistic description, including semantics, syntax and morphology. In this section we are concerned with various linguistic expressions of causation that are used in English texts.

Any causative construction involves two components, the cause and its effect. Let's consider the following example (Comrie 1981):

"The bus fails to turn up. **As a result**, I am late for a meeting".

In this example, 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 contain relevant keywords like *cause, effect, causation relationship*, etc, while implicit causative constructions are more complex, involving inference based on semantic analysis and background knowledge.

3.1 Explicit causative constructions

Linguists consider that explicit cause-effect lexico-syntactic patterns can be expressed in English texts in the following ways:

- Causal connectives
- Causation verbs
- Conditionals

- Causative adverbs and adjectives

3.1.1. Causal connectives. Using a compilation of other works (Greenbaum 1969), (Halliday and Hassan 1976), (Quirk et al. 1972), Altberg (Altberg 1984) classified the causal connectives into the following types:

A Adverbial causal link

B Prepositional causal link

C Subordination

D Clause-integrated link

A. Adverbial causal links (*for this reason, the result that*) are constructions that link two clauses in order to form a causation relationship. They can be of two types: *anaphoric link* and *cataphoric link*.

The anaphoric link references back to an element in the preceding discourse, as in the example:

"The meaning of a word can vary a great deal depending on the context. **For this reason**, pocket dictionaries have a very limited use."

The cataphoric link references ahead to an element in the preceding discourse. For instance,

"Labor government which came to power in 1996 stalled this process, **with the result that** Malta was not among the six countries opening membership negotiations with the EU in March 1998."

B. Prepositional causal links (**because of**, **thanks to**, **due to**) usually link a noun phrase with a clause, or two noun phrases in an apposition:

"A local man was kept off a recent flight **because of** a book he was carrying."

"Health problems, **due to** global warming, are predicted to increase at the end of the century".

C. Subordination causal links can be further classified in the following subcategories:

- **resultative conjunctions** (*because, as, since, for, so, so that*), e.g.:
  "The colonies came to realize they had to separate from England, **so** they started the Revolutionary War."
- **structural link** by a non-finite *ing*-clause, e.g.,
  "**Being** cloudy, the experiment was postponed."
- **correlative comparative construction**, e.g.,
  "The traffic was **so** heavy **that** I couldn't arrive on time."

D. The last causal link of this category is the **clause-integrated link** which is either part of the subject (called , *thematic link* (1)) or the predicate of the clause (*rhematic link* (2)):

(1) "The new satellite was named ASUKA (flying bird). **The reason was that** the satellite in orbit looks like a migratory bird soaring into deep space."

(2) "It is not a myth that world hunger **is due to** scarcity of food."

3.1.2. Causation verbs. Many linguists focused their attention on causative verbal constructions, mainly because their study involves the interaction of formal syntax and semantic analysis of the language.

The first important and accepted classification of causative verbs was done in 1969 by two Soviet linguists (Nedjalkov and Silnickij 1969). They proposed a lexical decomposition which tries to build a taxonomy of causative verbs according to whether they define only the causal link or the causal link plus other components of the two states of affairs that are causally related:

- **Simple causative** (**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*. For example,

  "Earthquakes **generate** tidal waves."
- **Resultative causatives** (**kill**, **melt**, **dry**, **break**, **drop**, etc.)
  These verbs refer to the causal link plus a part of the resulting situation.
- **Instrumental causatives** (**poison** (killing by poisoning), **hang**, **punch**, **clean**, etc.)
  These causatives express a part of the causing event as well as the result.

Based on this general characterization of verbal causative constructions, another linguist (Comrie 1973) introduced a new semantic classification that is now universally accepted by the linguistics community:

- **Analytic causatives** (periphrasic causatives)
- **Morphological causatives**
- **Lexical causatives**

An **analytical causative** describes a separate verb associated with the causative meaning. Usually it can be also used as a verb per se, with a different, but often related, meaning. For example, the verb *make*, can occur as the causative verb in a construction like *I made him do the homework*, or as a main verb in its own right, as in *I did the homework*. The most cited classes of analytical causatives, based on their fine semantic distinctions, are **make**, **get**, **have**:

"I **made** him do the homework."

"I **got** him to do the homework."

"I **had** him do the homework."

Morphological causatives are single words where the causative meaning is conveyed by a special morpheme or morphological process. In English, the suffixes $-en$ and $-ify$ form morphological causatives, such as *blacken, sweeten, thicken, nullify, liquefy, verify*, etc.

The **lexical causative** takes in words like *kill* and *feed*, which appear to be in a direct semantic relationship with other verbs, such as *die* and *eat*, but where the causative relation receives no formal expression at all.

The most common semantic distinction drawn from this classification (Goddard 1998) is that between direct and indirect causation. Wierzbicka (Wierzbicka 1980) showed that lexical causatives like *kill* or *break* imply more direct causation than the nearest analytic versions, such as *cause to die*, or *make it break*. Thus, according to Haiman (Haiman 1985), the greater the ***linguistic distance*** between the elements representing the cause and the effect, the greater the ***conceptual distance*** between them also.

3.1.3. Conditionals. Conditionals are typically expressed in English as sentences of the form "If S1, then S2", where S1 is the antecedent and S2 is the consequent. The state of affairs described in the antecedent is asserted to be a sufficient condition on the circumstance described in the consequent:

"`If` it rains, `then` I will stay at home."

However, conditionals in English may express something more than necessary and sufficient conditions: they may express causation (1), and hence, temporal succession (2).

(1) "`If` John studies, `then` he will pass the exam."

(2) "`If` John didn't pass, `then` he couldn't have studied."

Conditionals are complex linguistic structures, because they can also express inferences.

Dancygier (Dancygier 1993) offered a new and indepth analysis of English conditional sentences. She classifies conditional constructions according to time-reference and modality and shows how the basic meaning parameters of conditionality correlate to formal parameters of the linguistic constructions that are used to express them.

A theory of conditionals aims to give an account of the conditional construction which explains when conditional judgments are acceptable, which inferences involving conditionals are good inferences, and why this linguistic construction is so important. Despite intensive work of great ingenuity, this remains a highly controversial subject.

3.1.4. Causative adverbs and adjectives. Cresswell (Cresswell 1981) showed that some adverbs and adjectives have a causal element in their meaning. For example,

"Brutus `fatally` wounded Caesar."

"Caesar's wound was `fatal`."

Cresswell classified the causal adverbs in the following categories:

- Adverbs of perception (***audibly, visibly***)
- Adverbs marginally perceptual (***manifestly, patently, publicly, conspicuously***)
- Adverbs that involve the notion of a result whose properties are context dependent (***successfully, plausibly, conveniently, amusingly, pleasantly***)
- Adverbs that suggest tendencies, liabilities, disposition or potencies (***irrevocably, tenuously, precariously, rudely***)
- Adverbs that refer to effects (***obediently, gratefully, consequently, painfully***)
- Adverbs of means (***mechanically, magically***)

As it was showed in this classification, the English language provides many ways of expressing cause-effect relations. However, almost all these linguistic patterns are ambiguous, referring to a specific relation based on a particular context. Thus, any attempt of automatic detection and extraction of causation relations from text has to deal with the disambiguation of the corresponding causation constructions.

3.2 Implicit causative constructions

- Complex nominals
- Verbs of implicit causality
- Discourse structure

3.2.1. Complex nominals expressing causation. Complex nominals represent one of the most difficult problems in Natural Language Understanding, mainly because they require complex semantic analysis. Complex nominals are noun phrases formed as a succession of nouns and/or adjectives, like for instance "***English teacher, blueberry muffin, mortgage rate***", etc. What makes the analysis of this construction difficult is the ambiguity of the relation that exists between the underlying nouns. In general, a lexical unit is likely to prove ambiguous if it has more than one sense. Thus, the interpretation of the semantic structure is very difficult, due to the fact that no semantic relation between constituents is formally indicated. To be able to interpret them adequately, extralinguistic knowledge related to the semantic content of their components and the way they relate syntagmatically is required.

One of the relations that can link two nouns in a complex nominal is the causation relation. Complex nominals that express causation can be of the following types, in the decreasing order of ambiguity:

1 **Complex nominals with an implicit causation relationship.**
This is the most difficult type, as it requires inference based on semantic analysis and world knowledge.

- `NP1 NP2` ⇒ NP1 causes NP2 (e.g., *cold tremble*)
- `NP2 NP1` ⇒ NP1 causes NP2 (e.g., *malaria mosquitoes*)

2 **Complex nominals with an explicit but ambiguous causation relationship** (semi-explicit).
Here the relation is expressed by a verb, or a preposition, but which is highly ambiguous.

- `NP1-produced NP2`
- `NP1-inducing NP2`
- `NP1-provoking NP2` (e.g., "anxiety-`provoking` situations")

- `NP1-related NP2`
- `NP2 with NP1` (e.g., "gastroenteritis `with` Salmonella enteritis")
- `NP2 by NP1` (e.g., "infection by Staphylococcus aureus")

3 **Complex nominals with an explicit, unambiguous causation relationship.**
The relation is expressed explicitly, through the use of verbs like *cause*.

- `NP1-causing NP2` (e.g, "disease-`causing` bacteria")
- `NP2-caused NP1` (e.g., "infection-`caused` hoarness")

There are situations where the head noun in the complex nominal construction is modified by an adjective, usually derived itself from a noun (Levi 1979). For instance, *thermal stress* which comes from *heat stress*. The same typological analysis applies for this type of complex nominals also.

3.2.2. Implicit causality of verbs. Caramazza and his colleagues (Garvey and Caramazza 1974), (Caramazza et al. 1977) observed that in sentence fragments such as (1) and (2) readers prefer to interpret the pronoun as referring to one or other of the two potential referents, despite the lack of disambiguating gender information. The pronoun in (1) is preferentially interpreted as referring to the burglar, and in (2) it is preferentially interpreted as referring to the policeman.

(1) "The burglar `confessed` to the policeman because he was sorry for what he did."

(2) "The actor `admired` the policeman because he was brave."

Caramazza and others (Garham and Oakhill 1985) have argued that these preferences are due to a property of *implicit causality* that belongs to the verbs. Some verbs, including *confessed*, are said to impute implicit causality to the agent, meaning that in (1) the burglar is the instigator of events described in the sentence, and any continuation that attempts to explain the cause of these events is likely to be concerned with this character. In contrast, verbs such as *admired* are said to impute *implicit causality* to the patient of the sentence, so in (2) it is the policeman who causes the admiration event to occur. Therefore, the answer to the question *why did the actor admire the policeman?* is likely to be concerned with the policeman. Importantly, implicit causality only favors one of the participants as the referent of the pronoun.

3.2.3 Discourse structure. Discourse psychologists have developed and tested models that predict what inferences are generated on-line during comprehension. When reading a novel, for example, the following classes of knowledge-based inferences are potentially generated: goals and plans, character traits, characters' knowledge and beliefs, character emotions, causes of events, consequences of events and actions, spatial relationships among entities, etc.

Let's consider the sentences "The dragon was dragging off the girl. A hero came and fought the dragon" (Graesser et al. 1997). According to Graesser, Mills and Zwaan, there are five classes of inferences that might be encoded when the second sentence is read:

1 *Superordinate goal (motive).* The hero wanted to rescue the girl.

2 *Superordinate goal or action.* The hero threw a spear.

3 *Causal antecedent.* The girl was frightened.

4 *Causal consequence.* The hero married the girl.

5 *Static property.* The dragon has scales.

4. Semi-automatic detection of causation relationships

In this section we propose a method for automatic detection of causation patterns and semi-automatic validation of ambiguous verbal lexico-syntactic patterns that express causation.

The algorithm for the detection of lexico-syntactic patterns that refer to causation consists of three major procedures. The first procedure discovers lexico-syntactic patterns that can express the causation relation, the second procedure provides semantic constraints imposed by causation, and the third procedure validates and ranks the ambiguous patterns acquired based on semantic constraints on nouns and verbs.

4.1 Automatic discovery of lexico-syntactic patterns referring to causation

The causation relation can be expressed in text in various ways, from explicit to implicit, and from intra to extra-sentential patterns. One of the most frequent explicit intra-sentential pattern that can express causation is <*NP1* VERB *NP2*>.

In this paper we focus on explicit intra-sentential syntactic-patterns of the form <*NP1* VERB *NP2*>, 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 (Moldovan and Girju 2001):

Procedure 1. *Discovery of lexico-syntactic patterns:*

`Input:` semantic relation R

`Output:` list of lexico-syntactic patterns expressing R

1 Pick a semantic relation R (in this paper, CAUSATION)

2 Pick a pair of noun phrases C_i, C_j among which R holds.

In order to get as many causation patterns as possible, we repeated step 2 for a list of noun phrases extracted from WordNet 1.7. WordNet (Miller 1995) contains 17 semantic relations

(Harabagiu and Moldovan 1998): IS-A, reverse IS-A, MERONYMY/HOLONYMY, ENTAIL, CAUSE-TO, ATTRIBUTE, PERTAINYMY, ANTONYMY, SYNSET (SYNONYMY), etc. 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*". Given the fact that almost all these verbs have nominalizations, it is easy to find noun concepts among which the WordNet causation relations hold. 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%).

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. Retain only the sentences containing the pair. From these sentences, determine automatically all the patterns <*NP1* verb/verb_expression *NP2*>, where NP1 - NP2 is the pair considered.

The result is a list of verbs/verbal expressions that refer to causation. Some of these verbs are always referring to causation, but most of them are ambiguous, in the sense that they express a causation relation only in a particular context and only between specific pairs of nouns. For example, <*NP1 causes NP2*> refers always to causation, but this is not true for <*NP1 produces NP2*>. 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 (NP1 verb NP2). As is described in the next subsection, the relationships are disambiguated and ranked and only those referring to causation are retained.

4.2 Semantic constraints imposed by causation

Because the exact disambiguation of the verb sense is often very difficult, we try to validate the lexico-syntactic patterns using a coarse-grain approach. The procedure consists of detecting the constraints necessary and sufficient on nouns and verb for the pattern <*NP1* VERB *NP2*> such that the lexico-syntactic pattern indicates a causation relationship.

Semantic constraints on nouns NP_1 and NP_2

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.

In a relationship of the form <*NP1* VERB *NP2*>, the nouns NP1 (*cause_noun*) and NP2 (*effect_noun*), can express explicit or implicit state of affairs. The following four situations can occur:

1 *cause_noun* and *effect_noun* are explicit state of affairs.
e.g: "*Earthquakes* cause *tidal waves*".

2 *effect_noun* expresses an explicit state of affair, and *cause_noun* an implicit one.
e.g: "*John* caused the *disturbance*".

3 *cause_noun* shows an explicit state of affair, and *effect_noun* an implicit one.
e.g: "Sometimes *rain* can cause you *bad days*".

4 *cause_noun* and *effect_noun* are implicit state of affairs.
e.g: "*John* caused her really *bad days*".

Examples 2, 3 and 4 denote a causation relationship as the verb *cause* indicates, but the relation is not fully explicit. John cannot cause directly a psychological state (e.g., the *disturbance*), but the action John undertook caused it. In this paper we focus only on the situations 1 and 2, as they are the most frequently used in texts.

Given this approach the system selects automatically the causation classes with the following procedure:

Procedure 2.

Input: NP_1, NP_2, verb.

Output: Semantic constraints on nouns and verb.

STEP 1. **Semantic constraints on NP_2**
In step 2 of Procedure 1, the system detected automatically from WordNet a list of noun pairs NP_1 - NP_2 that are in a causation relationship. For each noun NP_2 occupying the EFFECT position in these relationships, select as causation class the most general subsumer in WordNet for that given sense. For example, the most general subsumer of the word *excitement* (#1/4) in WordNet is *psychological feature*. All the EFFECT nouns in the extracted causation pairs represent entities that express explicit state of affairs.

At the end of this step, the system detected the following causation classes: *human action, phenomenon, state, psychological feature*, and *event*. Our assumption is that these classes represent causation categories, and anything else that is not in this list refers to noncausation.

STEP 2. **Semantic constraints on NP_1**
We noticed from the corpus created in Procedure 1 that metonymy occurs with high frequency in causation relationships, but mostly on the CAUSE position, and quite rarely on the EFFECT position.

This observation is also supported by the large number of classes obtained for the NP_1 nouns on the *cause* position with the procedure described above. This shows that the CAUSE nouns can be represented by almost any noun. Thus, we use here only a soft constraint which would help validate the relationships in some special cases explained later in section 4:
soft constraint on CAUSE: the noun should have as subsumer the concept *causal agent* in WordNet. For example, the second most general subsumer of the word *drug* in WordNet is *causal agent*.

STEP 3. **Semantic constraints on verbs**
We ranked the verbs/verb expressions extracted in step 3 of Procedure 1 based on their ambiguity and frequency levels in WordNet. In WordNet, verbs are represented in *synsets*, which are lists of synonyms for that verb, and each verb can have multiple senses. For a given verb, in WordNet 1.7 the senses are ranked based on the number of times each sense occurs in the semantically tagged corpus used by the WordNet lexicographers. Based on the observation on WordNet of the extracted verbs, we considered the following categories of constraints along with their thresholds:

1 low ambiguity: - if the number of senses for the verb considered <= 7

2 high ambiguity: if the number of senses for the verb considered > 7

3 low frequency: - if (the frequency for that particular sense < the sum of the frequency of all other senses) or (the frequency for that particular sense < 30)

4 high frequency: if (the frequency for that particular sense >= the sum of the frequency of all other senses) or (the frequency for that particular sense >= 30)

Table 1 shows the verbs extracted with Procedure 1 ranked according with the constraints defined above.

For example, the verb *make* is ranked at the end because it is highly ambiguous (there are 49 senses in WordNet 1.7 for this verb) and occurs with high frequency (for sense #5/49 (*cause*) there are 79 occurrences in the WordNet tagged corpus). Thus, the sentence "*Greenspan* makes *a recession*" is highly ambiguous as it can be interpreted in two ways: either (1) as a causation relation if *recession* has the sense #1/4 ("*the state of the economy declines*"), or (2) as noncausative relation if *recession* has the sense #2/4 ("*a small concavity*").

Using the verb constraints presented in this step, the system ranks automatically the causation verbs discovered in step 3 of Procedure 1 in four classes based on two parameters: ambiguity and frequency (Table 1). The higher the ambiguity and frequency of the verb, the less changes there are for it to express causality.

4.3 Validation of causation patterns and ranking of causation relationships

Causation relations of the type NP_1 CAUSE_VB NP_2 can have different levels of ambiguity, based on the ambiguity information derived for each component word or expression. The more ambiguous the constituents, the more difficult, and thus unlikely it is to classify the relationship as causation. In this subsection we propose a five-level ranking of causation relationships, in the increasing order of ambiguity.

The algorithm for the validation and ranking of the causation relationships is an iterative procedure in which a step is followed if the condition in the previous step was not satisfied.
In this algorithm we consider as NP_1 and NP_2 only the head noun of the noun phrases extracted as cause and effect, as it occurs in WordNet (e.g, for the noun phrase "giant tidal wave", *tidal wave* is automatically selected).

Procedure 3.

`Input:` Ambiguous causation patterns

`Output:` Ranked list of causation patterns

Step 1.
If the EFFECT and CAUSE head nouns are monosemous and they belong to one of the causation classes, or are polisemous and all their senses belong to the causation classes, then classify the relationship as causation of *rank 1*.

For example, "*Hitler's invasion of Poland* provoked *the Second World War*".

Here, both *invasion* and *Second World War* have all their senses in causation classes, so even if the verb *provoke* is ambiguous, the relationship is detected as causation.

Step 2.
If the EFFECT head noun is monosemous and belongs to one of the causation classes, or is polisemous and all its senses belong to the causation classes, then classify the relationship as causation of *rank 2*. The rationale is that most of the time, the causation relations that deal with metonymy, have it expressed in the CAUSE position.

For example, "In 1958, it was *Bleustein-Blanchet* who sparked *a controversy* when he opened Le Drugstore, the American-inspired combination pharmacy, all-hours restaurant and gift store that now has branches at both ends of the avenue".
Here, despite the fact that the CAUSE is a metonymy (e.g., *Bleustein-Blanchet*), the causation relation is obvious as *controversy* is monosemous and its sense has the semantic class *human action*.

Low ambiguity High frequency	Low ambiguity Low frequency	High ambiguity Low frequency	High ambiguity High frequency
give rise (to) (#1/1)	stir up (#2,4/4)	create (#1/6)	start (#2/13)
induce (#1/4)	entail (#1/2)	launch (#7/8)	make (#5/49)
produce (#3/7)	contribute (to) (#3/4)	develop (#5/21)	begin (#3/10)
generate (#1/4)	set up (#8/15)	bring (#3/11)	rise (#7/19)
effect (#1/2, #2/2)	trigger off (#1/1)		
bring about (#1-4/5)	commence (#2/3)		
provoke (#1-2/4)	set off (#1,7/7)		
arouse (#1,4,5/7)	set in motion (#1/1)		
elicit (#1/3)	bring on (#1/3)		
lead (to) (#2,3,5/15)	conduce (to) (#1/1)		
trigger (#1/2)	educe (#2/2)		
derive (from) (#2/5)	originate in (#1/1)		
associate (with) (#1/3)	lead off (#2/2)		
relate (to) (#1/6)	spark (#1/2)		
link (to) (#1/4)	spark off (#1/1)		
stem (from) (#1/2)	evoke (#1,2/4)		
originate (#1/3)	link up (#3/3)		
bring forth (#1/4)	implicate (in) (#2/2)		
lead up (#1/1)	activate (#1/5)		
trigger off (#1/1)	actuate (#1/1)		
bring on (#1/3)	kindle (#3/3)		
result (from) (#1,2/2)	fire up (#1/2)		
	stimulate (#3/7)		
	call forth (#1,2/2)		
	unleash (#1/3)		
	effectuate (#1/1)		
	kick up (#2/2)		
	give birth (to) (#2/2)		
	call down (#1/1)		
	put forward (#3/5)		

Table 1: Ambiguous causation verbs ranked based on ambiguity and frequency. The ambiguity increases from the left most column to the right.

Step 3.
If the EFFECT is represented by an enumeration of noun phrases and the head noun of at least one of them has all the senses in one of the causation classes, than the others also refer to causation in that context. Classify the relationship as causation of *rank 3*.

For example, in the sentence "*Fed* will **induce** *a recession and unemployment*" the effect *unemployment* is monosemous and belongs to the causation class *state*. Thus, the effect noun *recession* is disambiguated and its interpretation as sense #2 (*niche, corner*) is eliminated.

Step 4.
If the noun phrase representing the EFFECT is ambiguous (at least one of its senses does not belong to a causation class) and the CAUSE respects the soft constraint defined in the previous subsection, then classify the relationship as causation of *rank 4*.

For example, in the sentence "*The drugs* **induce** *the growth of muscle tones*", the head noun *growth* has two senses (#4/7 and #7/7) that are in two noncausation classes (e.g., {*group, grouping*}, and respectively {*entity*}). In this case, the noun *drugs* disambiguates the relationship as it is monosemous and has *causal agent* as one of its hypernyms.

Step 5.
At this point, the remaining nouns representing the CAUSE and EFFECT are ambiguous and the only possibility of disambiguation comes from the restrictions imposed on the verbs.

For example, in the sentence "*The issue* **gives rise** *to a big concern*", both the CAUSE and the EFFECT are ambiguous. The noun *issue* can be "*an important question that is in dispute and must be settled*" (*psychological feature*, cf. WordNet), or "*one of a series published periodically*" (entity, cf. WordNet). The noun *concern* can refer to "*an anxious feeling*" (*psychological feature*, cf. WordNet), or "*commercial or industrial enterprise*" (*group, grouping*). In this case the relationship is considered causation only because the verb *give rise* is one of the less ambiguous and highly frequent verbs considered.

For all the remaining relationships, classify them based on the verbs' ranking shown in Figure 1.

5. Results

In this section we show the results obtained by the validation and ranking algorithm. For this experiment we used the TREC-9 (TREC-9 2000) collection of texts which contains 3GB of news articles from *Wall Street Journal, Financial Times, Financial Report*, etc. Using the causation verbs obtained in step 3 of Procedure 1, the system formed queries and searched the TREC collection. This way, for each verb there were selected 50 sentences that contained it. The new corpus thus formed (3,000 sentences) was part-of-speech tagged and parsed. For each head of the noun phrases in the CAUSE and EFFECT positions, the system determined automatically the most general subsumer for each sense. The algorithm presented in subsection 4.3 was implemented and the system gave as output 1,321 causation relationships $<NP_1$ verb $NP_2>$, ranked by generality.

The results were validated by comparison with human annotation. We asked two subjects, other than the authors, to rank a list of 300 relationships from which only 230 were referring to causation, as detected by our algorithm. Out of the 300 relationships the subjects selected as causation relationships only 151 on average (Table 2). In what concerns the rating of the causation relationships, it differed from one subject to another with about 36%, and from the system's output by 48%.

The accuracy obtained by our system in comparison with the average of two human annotations was 65.6%.

	System	Human annotator 1	Human annotator 2
Rank 1	37	30	21
Rank 2	73	58	43
Rank 3	28	18	22
Rank 4	92	56	54
Total	230	162 (70.43%)	140 (60.87%)

Table 2: Comparison with human annotation and accuracy obtained for the 230 causation relationships (the percentages in parentheses represent the accuracy obtained by the system reported to the human annotator).

6. Cause-Effect Questions

Causation relationships occur in text with high frequency, 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 way, an in depth semantic analysis of cause-effect relations requires a ranking of causation patterns. This ranking proves to be very useful for applications like Question Answering.

Causation questions can be roughly classified in the following 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, causal relation*, etc.)

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

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

"Which were the **consequences** of Mt. Saint Elena erruption 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, bring about, 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?"

"Why are underwater volcanoes **dangerous** to fish?"

3 **Implicit causation questions**
This type of questions involves reasoning, based on deep semantic analysis and background knowledge. This questions are usually introduced by the semantic types *why* and *what*.

This type of causation questions can be further classified in two important subtypes:

- *Causation questions disambiguated based on the semantic analysis of question keywords*

  "**Why** did Socrates **die**?"

  "**What killed** Socrates?"

  "Do volcanic eruptions **serve** any useful **purpose**?"

  It is recognized that questions of type *what*, and even *why*, are ambiguous, and usually the question is disambiguated by other keywords in the question. The verb *kill* is a causation verb meaning *cause to die*, so the second question asks for the cause of the Socrates' death.

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

  Nikko Tinbergen, one of the 'fathers' of ethology, argued that there are really only four basic questions that can describe the causes of behavior. He showed that for the *why*-type questions there are four different possibilities of answers:

  Question: *Why do robins sing in the spring?*

  - **Causation.** (What is the cause?)
    Answer: "because increasing day length stimulates hormonal action".
  - **Development.** (How does it develop?)
    Answer: "Males learn their behavior from their father and neighbors" (because they are adult males).
  - **Origin.** (How did it evolve?)
    Answer: "Song evolved as a means of communication early in the avian lineage".

- **Function.** (What is the function?)
  Answer: "to attract mates" (survival)

The algorithm for automatic extraction of causation relations presented in section 4 was tested on a list of 50 natural language questions using a state-of-the-art Question Answering system (Harabagiu et al. 2001). The list of questions were representative for the first two categories of causation questions presented above, namely *explicit* and *ambiguous causation questions*. Figure 3 shows two examples of questions from each class. The questions were tested on the QA system with and without the causation module included.

7. Discussion and Conclusions

The approach presented in this paper for the detection and validation of causation patterns is a novel one. Even if the method is not fully automated, it brings considerable improvement in time and user work compared with other previous attempts (Garcia 1997), (Khoo et al. 2000). Khoo at al. obtained a better accuracy, but they restricted their text corpus to a medical 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 patterns application on text.

We present here a few problems that have caught our attention.

1 The data is sparse. For some of the verbs considered, a large number of occurrences referred to other senses than causation.

2 For this experiment we did not consider semantic constraints among nouns and verb. This way, the system did not filter out some exceptions like *The plant generated energy*. Even if the EFFECT *energy* does not have all senses in the causation classes, the relationship was selected by the system because the verb *generate* was ranked on top of the verbs' list (low ambiguity, high frequency). However, this is not a causation relationship, as here the verb *generate* means "*produce energy*".

We intend to extend the analysis to other *causation* patterns and devise a general algorithm for the detection and especially for the validation of causation patterns.

Acknowledgements. This research was partially funded by the AQUAINT Program, ARDA.

References

B. Altberg. Causal linking in spoken and written English. In *Studia Linguistica*, 38(1), 1984.

A. Caramazza, E. Grober, C. Garvey, and J. Yates. Comprehension of anaphoric pronouns In *Journal of Verbal Learning and Verbal Behavior*, 1977.

Bernard Comrie. Aspect. Cambridge University Press, Cambridge, UK, 1976.

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

Max Cresswell. Adverbs of Causation. In *Words, Worlds, and Contexts: New Approaches in Word Semantics*, Eikmeyer, Rieser (eds), 1981.

Barbara Dancygier. Interpreting conditionals. In *Journal of Pragmatics*, 19: 435-452, 1993.

Sanda Harabagiu and Dan Moldovan. Knowledge Processing on Extended WordNet. In *WordNet: An Electronic Lexical Database and Some of its Applications*, Editor C. Fellbaum, MIT Press, 1998, pages 379-405.

M. Halliday and R. Hasan. Cohesion In English. London, Longman, 1976.

Sanda Harabagiu, Dan Moldovan, Marius Pasca, Rada Mihalcea, Mihai Surdeanu, Razvan Bunescu, Roxana Girju, Vasile Rus and Paul Morarescu. The Role of Lexico-Semantic Feedbacks in Open-Domain Textual Question Answering. In the *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics (ACL-2001)*, Toulouse, France, July 2001.

Arthur Graesser, Keith Millis, and Rolf Zwaan. Discourse Comprehension. In *Annual Review of Psychology*, 48, 163-189, 1997.

S. Greenbaum. Studies in English Adverbial Usage. London: Longman, 1969.

D. Garcia. COATIS, an NLP system to locate expressions of actions connected by causality links. In *Knowledge Acquisition, Modeling and Mangement*, Proceedings of the Tenth European Workshop, EKAW '97, pages 347-352.

A. Garham and J. Oakhill. On-line resolution of anaphoric pronouns: Effecs of inference making and verb semantics. In the *British Journal of Psychology*, 1985.

Garvey and Caramazza. Implicit causality in verbs. In *Linguistic Inquiry*, 5:549-564, 1974.

Cliff Goddard. Causatives. In *Semantic Analysis: A Practical Introduction* University of New England, Australia (Oxford Textbooks in Linguistics), 1998.

John Haiman. Natural syntax. Cambridge: Cambridge University Press, 1985.

Jerry Hobbs. Causality. In the Proceedings of the *Fifth Symposium of Logical Formalizations of Commonsense Reasoning*, New York, NY, May 2001.

L. Joskowiscz, T. Ksiezyk and R. Grishman. Deep domain models for discourse anaysis. In *The Annual AI Systems in Government Conference*, Silver Spring MD, pages 195-200.

R.M. Kaplan, and G. Berry-Rogghe. Knowledge-based acquisition of causal relationships in text. In *Knowledge Acquisition*, 3(3), 317-337.

Christopher Khoo, Syin Chan and Yun Niu. Extracting Causal Knowledge from a Medical Database Using Graphical Patterns In *Proceedings of 38th Annual Meeting of the ACL*, Hong Kong, 2000, pages 336-343.

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

Judith Levi. The Syntax and Semantics of Complex Nominals. New York: Academic Press, 1979.

Question Class	Question	Answer	
		QA without causation module	QA with causation module
Explicit Causation Questions	What causes post-traumatic stress disorder?	Post-traumatic Stress Disorder - What are the Symptoms, Causes and Treatments?	Post-traumatic stress disorder is a condition resulting from a traumatic event outside the range of a person's normal experience.
	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 and to the deterioration and soiling of stone and paint on buildings, statues, and other structures of cultural significance.
Ambiguous Causation Questions	What can trigger an allergic reaction?	The molecular weight of the protein is also consistent with something that can trigger an allergic reaction	An antigen producing an allergic reaction. is defined as an allergen.
	What phenomenon is associated with volcanoes?	.. that deglaciation are actually associated with increased volcanic activity..	There are often earthquakes generated by volcanism..

Table 3: Examples of cause-effect questions tested on a Question Answering system.

Jong Gyun Lim. Planning in AI and Text Planning in Natural Language Generation. Technical Report (CUCS-038-92), Columbia University, NY, 1992.

G.A. Miller. WordNet: A Lexical Database. *Communication of the ACM*, vol 38: No11, pages 39–41, 1995.

Dan Moldovan and Roxana Girju. An Interactive Tool For The Rapid Development of Knowledge Bases. In the *International Journal on Artificial Intelligence Tools (IJAIT)*, vol 10., no. 1-2, March 2001.

V.P. Nedjalkov and G. Silnickij. The topology of causative constructions. In *Folia Linguistica* (6), 1969, pages 273-290 (german translation)

Charles Ortiz. A Commonsense Language for Reasoning about Causation and Rational Action. In *Artificial Intelligence*, Vol. 111, No. 2, 1999.

Judea Pearl. Causality. Cambridge University Press, 2000.

R. Quirk, S. Leech, S. Greenbaum, G. Leech, and J. Svartvik. A grammar of Contemporary English. London, Longman, 1972.

John F. Sowa. Knowledge Representation: Logical, Philosophical, and Computational Foundations, Brooks/Cole, Pacific Grove, CA, 2000.

Text REtrieval Conference. http://trec.nist.gov 2000

Leonard Talmy. Semantic Causative Types. In *Syntax and Semantics*, vol. 6. The Grammar of Causative Constructions, ed. Masayoshi Shibatani, Academic Press, 1976.

Anna Wierzbicka. Causatives Re-visited. In *The New Psychology of Language*, London: Routledge, 1980.

Robert W. Wilenski. Understanding Goal-based Stories. Ph.D. thesis, Yale, 1978.