From *Text-Based Intelligent Systems*, edited by Paul Jacobs, Lawrence Erlbaum Associates, 1992.

Direction-Based Text Interpretation as an Information Access Refinement

Marti A. Hearst
Computer Science Division, 571 Evans Hall
University of California, Berkeley
Berkeley, CA 94720
and
Xerox Palo Alto Research Center
marti@cs.berkeley.edu

Abstract

A Text-Based Intelligent System should provide more in-depth information about the contents of its corpus than does a standard information retrieval system, while at the same time avoiding the complexity and resource-consuming behavior of detailed text understanders. Instead of focusing on discovering documents that pertain to some *topic* of interest to the user, an approach is introduced based on the criterion of *directionality* (e.g., Is the agent in favor of, neutral, or opposed to the event?). A method is described for coercing sentence meanings into a metaphoric model such that the only semantic interpretation needed in order to determine the directionality of a sentence is done with respect to the model. This interpretation method is designed to be an integrated component of a hybrid information access system.

1 Introduction

In the light of the increasing availability of computer-accessible full text, an important goal of a Text-Based Intelligent System is to provide a means for answering questions about documents' contents. Standard information retrieval systems have sophisticated methods for grouping documents according to similarities among their terms and the terms in user queries (see [Salton, 1988], Croft & Turtle, this volume, for overviews). Often the data for these systems are titles and abstracts, as opposed to full text documents. Using this data, given a document whose topic is known, other documents whose terms are similar to it are considered likely to be about the same topic. However, document similarity is only one of many useful criteria for accessing information, and the availability of full text opens up exploration of others.

One way of structuring a corpus is to sort the documents into categories based on their topical content. Current text systems accomplish this task with varying degrees of sophistication. For example, RUBRIC [McCune et al., 1985] allows the user to define an elaborate conceptual hierarchy, bottoming out on keywords, that classifies documents according to what topics they contain. (Hayes, this volume, describes related systems.) Assuming, then, that the main topic of a document can be determined, how can the document be further distinguished from others describing the same topic?

One way to distinguish a document from its neighbors is to answer specific questions about its contents. A good number of systems have been developed that look for answers to a set of predefined questions about a specific topic domain, notably those demonstrated at the Message Understanding Conferences [Lehnert and Sundheim, 1991]. Most of these systems require large amounts of task-specific domain knowledge and complex inferencing capabilities. The process of building up and representing the necessary knowledge bases is time-consuming and good coverage is difficult to achieve. For this reason, our question should be revised to: how can a document be distinguished from others describing the same topic without the costs associated with domain-dependent approaches?

What is needed is a classification criterion that applies to a wide range of text; a useful question relatively independent of domain. One such criterion is: where, according to the text, does a semantic attribute lie along a continuum between extremes? For example, given a corpus of newspaper articles and the topic "Environmental issues pertaining to wildlife refuges," one can inquire as to whether public figures are stated as being opposed to, neutral, or in favor of a proposed cleanup plan. More generally, articles can be classified according to how they answer the query "Is agent A in favor of event E?" Other examples of queries within this genre are: "Is situation S improving or worsening?" and "Is agent A1 dominating or being dominated by agent A2?" I call this criterion *directionality*, in contrast to the *topicality* criterion mentioned above. Directional queries are domain independent; whether the domain is wildlife refuges, mideast peace agreements, or urban policy, the question can still be applicable.

A mechanism that can classify an article based on the directionality criterion provides a precise interpretation of a narrow slice of the semantic content of the document. The goal is to avoid the expense of a full semantic analysis by *restricting* the type of information extracted from the text. How is sentence directionality to be determined? Clearly keyword-based analysis alone is not sufficient. Consider the classification criterion "Is the agent in favor of the event?" applied to the following pair of sentences:

(1a) The congresspersons introduced legislation to *lift* the ban on wastewater dumping.

(1b) The congresspersons introduced legislation to *support* the ban on wastewater dumping.

A difference of one word manages to reverse the attitude of the agents toward the situation, even though *lift* and *support* are not antonyms when out of context. To correctly distinguish these sentences, at least a partial understanding must come into play. However, the semantics need not be comprehensive – the interpretation mechanism can take advantage of the restricted nature of the query in order to minimize the degree of inference needed.

To this end I propose a sentence interpretation model called direction-based text interpretation (DTI). In direction-based text interpretation, isolated portions of a text are interpreted within the framework of a general, domain independent metaphoric model. This model is derived from Talmy's theory of force dynamics [Talmy, 1985] and involves coercing the meanings of sentences that satisfy the directionality criterion into a general conceptual framework and then interpreting the sentences with respect to that framework. When working within this restricted model, lexical items require assignment of a value to only one semantic attribute, thus circumventing the need for the large, complex knowledge bases required by full text understanding systems. Integrating this method with an information retrieval system should yield an incremental improvement in the text classification task.

In other words, DTI involves interpreting text in terms of a simple semantic model, only to the amount of detail necessary to accomplish the target task (to answer the query of interest). A sketch of the overall procedure is: relevant documents are selected by the system's information retrieval component, which makes use of domain-dependent keywords and phrases (assumed to already be supplied) that identify the target concepts (e.g., the system knows about lexical items involved in expressing a topic such as "wastewater dumping"). This information is used to isolate sentences that are likely to contain the answer to the target query (e.g., a sentence that refers both to "congresspersons" and "wastewater dumping" is a good candidate). Once a candidate sentence is found, it is parsed into a feature-structure form. As the analysis proceeds, pieces of the directional model are instantiated and linked together corresponding to elements of the parse. The resulting structure is interpreted in terms of the directional model, and the query is answered.

In summary, this chapter describes a domain-independent, coarse-level text interpretation method intended to be integrated into an intelligent information access system.

The next section, Section 2, explains the conceptual model which underlies direction-based text interpretation, and Section 3 describes mechanisms for interpreting sentences within this model.

Section 4 is a short discussion of related work, followed by, in Section 5, an outline of the retrieval paradigm into which direction-based text interpretation might be placed. Section 6 concludes the chapter with some questions about the feasibility of the approach.

2 The Conceptual Model: Applied Cognitive Linguistics

This section explains the conceptual model which underlies direction-based text interpretation. The *path model*, as I call it, is an extension of Talmy's force dynamic model with some influence ¿from Reddy's description of the conduit metaphor.

2.1 The Force Dynamic Model

Studies in cognitive linguistics have shown that in some cases a multitude of linguistic phenomena can be well described in terms of a general conceptual framework. Reddy [Reddy, 1979] describes how the *conduit* metaphor can be seen as underlying many English expressions about communication. In this framework a thought is schematized as an object which is placed by the speaker into a container that is sent along a conduit. The receiver at the other end is the listener, who then removes the objectified thought from the container and thus possesses it. Inferences that can be made about conduits (e.g., they can be blocked up, become full, etc.) are applied to notions of communication. For example, English speakers make statements such as "Your meaning did not come through," "I can't put this thought into words," and "She's sending me some kind of message with that remark." Thus it may be the case that speakers unconsciously structure their talk about communication within the framework of this metaphor.

Another example of a claim that an underlying conceptual framework is revealed by its expression in language is found in Talmy's theory of *force dynamics* [Talmy, 1985]. This theory describes how the interaction of agents with respect to force is lexically and grammatically expressed (focusing on English). Notions within the scope of force dynamics include: exertion of force by an agent, resistance to this force, overcoming this resistance, and so on. Talmy claims that force dynamics is a conceptual organizing system, one of the fundamental categories that structures and organizes meaning, akin to more familiar linguistic categories like number and aspect.

The force dynamic model posits the participation of two opposing entities, named the Agonist and the Antagonist. Each entity expresses an intrinsic force, tending either toward motion or toward rest. The relative strengths of the tendencies of the agents is important since the stronger entity is able to manifest its tendency at the expense of its opposer. The balance of relative strengths determines the interaction's resulting state.

To clarify these ideas, consider the following sentences (taken from [Talmy, 1985]):

(2a) The ball kept rolling because of the wind blowing on it.

(2b) The shed kept standing despite the gale wind blowing against it.

(2c) The ball kept rolling despite the stiff grass.

(2d) The log kept lying on the incline because of the ridge there.

In sentence (2a) the ball is seen as the Agonist and the wind as the Antagonist. The Agonist's tendency is toward rest but it is forced against its tendency by the stronger force of the wind. In sentence (2b) the Agonist (the shed) again has a tendency toward rest but in this case it is able to maintain this tendency against the opposing force of the Antagonist (the wind). In (2c) the Agonist is the ball tending toward motion. In this case the Antagonist's force does not succeed in reversing the Agonist's tendency, but in sentence (2d) the Antagonist (the ridge) overcomes the Agonist's tendency toward motion.

These examples illustrate only the simplest form of interactions. Talmy describes more complex examples in which the force interaction changes dynamically, and situations in which the Antagonist is the stronger of the two entities, but it remains "out of the way" of the Agonist, thus allowing the description of the concept "letting" (e.g., "The plug's staying loose let the water drain from the tank."). Talmy speculates that the traditional understanding of the scope of causation is too narrow, and should be expanded to include notions like letting since this evidence indicates that both letting and causing are expressed through the same conceptual framework.

Talmy further shows that certain force dynamic concepts have grammatical representation. For example, when the Agonist appears as the subject, the role of a stronger Antagonist can be expressed by conjoining a clause headed by "because" (as in (2a)). Similarly the interaction with a weaker Antagonist can be expressed by a clause headed by "although" or a prepositional phrase headed by "despite." The preposition "against" indicates force dynamical opposition as well as the particles "still" and "on" (e.g. "The ball was still rolling."). Talmy argues that the form "keep" (e.g., "The ball kept rolling.") might be considered an honorary auxiliary in the way that "have to" can act like a modal.

Whether or not the force dynamical system truly underlies the language user's conceptual system, it is a useful device for interpreting expressions of causal interaction. For example, the theory could be helpful for the problem of determining, given a sentence describing two entities engaged in an interaction, which one is relatively stronger.

2.2 The Path Model

If an agent favors an entity or event, the agent can be said to desire the existence or "well-being" of that entity. Furthermore, if the agent favors an impedance to the existence or well-being of an entity, then it can be said to be opposed to the entity. A useful heuristic can be inferred from this: namely, if an agent favors an entity's triumph in a force-dynamic interaction, then the agent favors that entity.

In a particular sentential description there may be a string of occurrences that affect the state of the entity. The force dynamic model does not have the expressive power to represent this, so it must be augmented. Instead of focusing on the relative strength of two interacting entities, the model should represent what happens to a single entity through the course of its encounters with other entities. Thus the entity can be schematized as if it were moving along a path toward some destination or goal. The entity may encounter barriers in its path, indicating that its tendency is being blocked. Agents independent of the entity have the power to introduce barriers, remove barriers, reinforce or weaken barriers, initiate the entity's journey, speed up or slow down the journey, or bring the entity to its destination. An agent's attitude toward an entity can be determined by how it chooses to affect the movement of that entity along its metaphorical path.

Empirical analysis of directional sentences reveals that the inferences that can be generated based on this path metaphor suffice to answer the directional query: "Is the agent in favor of the event?" This leads me to adopt the path model as the conceptual model for direction-based text interpretation.

Both Talmy and Reddy consider the base metaphors that they investigate to be at least part of the underlying meaning of some subset of linguistic utterances. However, in DTI the base metaphor is used as a *lingua franca* into which the meanings of *all* candidate sentences are coerced. This is useful for two reasons: first, once the system has a representation of the sentence based on the path model, it need perform only a restricted set of inferences. Second, since the model being mapped into is small (compared with mapping into a network of "real-world" knowledge), the assignment of semantic attributes to lexical items is simplified considerably.

Preliminary work reveals that the path model, with some minor modifications, can be applied to answer another general query, namely "Does the event E improve the situation S?" This includes subquestions such as "Does the drug cure the disease?" and "Is the financial situation improving?" These queries all have a directional component.

3 Determining Directionality via the Path Model

A descriptive theory like force dynamics is a tool for describing how a conceptual framework is expressed in language, rather than prescribing how to interpret sentence meanings. This subsection presents an initial description of how sentences can be interpreted in terms of the path model.

However, before launching into the details of interpretation, we should consider what kinds of results the system is expected to produce. Consider again sentence (1a). Assume the user is interested in the opinions of policymakers towards issues involving wastewater management, and that a mechanism for recognizing simple noun phrases involving policymakers and wastewater management is in place. The goal upon encountering this sentence is to determine that in this case the policymakers are "the congresspersons," the target entity is "wastewater dumping" and the direction is "pro," i.e., the congresspersons favor the wastewater dumping. It is possible also to draw two other directional conclusions from this example, namely: (a) the congresspersons favor legislation and (b) the congresspersons are opposed to the ban on wastewater dumping. Conclusion (a) in this case is not very interesting because it is quite general; it indicates a mechanism which enables the desired result. If the information were more specific, as in "The congresspersons introduced bill number AJ23 ..." then this information might be worth reporting. Conclusion (b) is undesirable because it can be analyzed more thoroughly, presenting the user with a more concise result. Both (a) and (b) can be concluded using the mechanism described here; the end application should indicate whether or not partial results like these be reported.

3.1 Path Actions

Although Talmy's description indicates some grammatical patterns that are involved in the expression of the force dynamical model, it relies as well on the meanings of the open-class lexical items. For example, in sentences (2a-d) the reader must know that "rolling" indicates a tendency toward motion and "standing" indicates a tendency toward rest. Since one of the main reasons for using the path model as the basis of analysis is to make the interpretive process simple, it is desirable to avoid an open-ended semantics-assignment task.

To this end, I define a set of *path actions* that represent the semantic components of the path model, i.e., they represent what kinds of actions can take place within the model. These actions are: *ENABLE*, *BLOCK*, *REMOVE-BLOCK*, *ACCELERATE*, *DECELERATE*, and *NEUTRAL*. Their semantics are glossed as:

> *ENABLE:* allow or help the entity to move along a path toward its destination
>
> *BLOCK*: impose a barrier in the entity's path toward its destination
>
> *REMOVE-BLOCK*: remove a barrier from the entity's path
>
> *ACCELERATE*: increase or intensify the potency of the entity's current path-movement tendency; speed up
>
> *DECELERATE*: decrease in potency the entity's current path-movement tendency; slow down
>
> *NEUTRAL*: no effect

Other actions, such as *MISDIRECT*, could have been included but those listed have been found to be sufficient for the sentences examined.

As noted above, in the path model it is not unusual for an entity's progress to be affected by a series of actions. Therefore, I define a function that takes as arguments two path actions and returns a path action; Table 1 shows the results of all possible function applications. For example, a *REMOVE-BLOCK* action applied to a *BLOCK* yields an *ENABLE*. For simplicity, I will represent this function as a binary (right associative) operator, indicated by $\triangleright$.

When answering the target query, the directionality of two separate components – the attitude of the agent and the progress of the entity or event – must be determined. There are expressions

▷	*NEUTRAL*	*BLOCK*	*REM-BLOCK*	*ENABLE*	*ACCEL.*	*DECEL.*
BLOCK	*BLOCK*	*ENABLE*	*ENABLE*	*BLOCK*	*BLOCK*	*ENABLE*
REM-BLOCK	*REM-BLOCK*	*ENABLE*	−−	−−	−−	*ENABLE*
ENABLE	*ENABLE*	*BLOCK*	*REM-BLOCK*	*ENABLE*	*ACCEL.*	*DECEL.*
ACCEL.	*ACCEL.*	*BLOCK*	*REM-BLOCK*	*ACCEL.*	*ACCEL.*	*DECEL.*
DECEL.	*DECEL.*	*ENABLE*	*ENABLE*	*BLOCK*	*DECEL.*	*ACCEL.*

action in row is applied to action in column
NEUTRAL cannot appear as the first argument

Table 1: Results of Applying One Path-Action to Another

that indicate the agent's attitude relatively directly (e.g., "favored," "denounced," and "lauded"). The attitude of the agent can be expressed indirectly, however, as seen in (1a), where the legislator's legislation introduction action is taken to indicate that the legislator favors the legislation. If a series of action applications ends in a state which is one of *ENABLE or ACCELERATE,* then the agent is said to favor the entity that the action is applied to. If the action sequence ends up as a *BLOCK* or a *DECELERATE*, then the agent is assumed to oppose the target entity. Otherwise, no opinion is expressed or can be determined.

The lexical items that comprise a sentence constituent determine which actions are associated with the constituent. For example, from sentence (1a), the simple noun phrase "the ban" is assigned the action *BLOCK*, and the verb group "to lift" is assigned the action *REMOVE-BLOCK*. A constituent whose path action is *NEUTRAL* acts as an end point in a chain of action applications. When this occurs, the lexical items that make up the *NEUTRAL* constituent are considered to be the effected entity, and are placed in a predicate that represents the results of the series of path actions. For example, if "wastewater dumping" is *NEUTRAL*, then "to lift the ban on wastewater dumping" induces the following sequence of operator applications:

REMOVE-BLOCK ▷ *BLOCK* ▷ *NEUTRAL* ⇒
REMOVE-BLOCK ▷ *BLOCK* ⇒
ENABLE ⇒
favor(A, "wastewater dumping")

(The term *A* represents the agent whose attitude is under scrutiny. The example fragment "to lift the ban on wastewater dumping" does not specify an agent, so *A* is left unbound.)

If sentence (1b) is processed instead, "to support" is assigned the *ACCELERATE* action and the sequence would be:

ACCELERATE ▷ *BLOCK* ▷ *NEUTRAL* ⇒
ACCELERATE ▷ *BLOCK* ⇒
BLOCK ⇒
oppose(A, "wastewater dumping")

The *BLOCK* is "sped up" or "strengthened," rather than removed.

The path actions interact in some interesting ways, as shown in Table 1. Notice that the "negative" types, *BLOCK* and *DECELERATE*, tend to flip the polarity of the action they are applied to, while the "positive" types, *ENABLE* and *ACCELERATE*, leave the polarity unchanged. Comparing the phrases "increased the restrictions" and "reduced the restrictions," we see that the lexical item "increased" indicates the *ACCELERATE* action, "reduced" the *DECELERATE* action, and "the restrictions" indicates a *BLOCK*. Analyzing "increased the restrictions" produces:

$$\textit{ACCELERATE} \;\triangleright\; \textit{BLOCK} \quad \Rightarrow \textit{BLOCK}$$

whereas "reduced the restrictions" produces:

$$\textit{DECELERATE} \;\triangleright\; \textit{BLOCK} \quad \Rightarrow \textit{ENABLE}$$

The heuristic that motivates these transformations is: in the political arena, if an agent is seen as favoring the reduction in potency of a barrier, then in actuality that agent wants the barrier removed entirely but is supporting the reduction as a compromise. Thus to favor a *DECELERATE* on a *BLOCK* is to desire an *ENABLE*. Similarly, "limits," "restrictions," and "ceilings" are often proposed as compromised alternatives to outright stoppages, and in this model they are cast as *BLOCK*s. If the system produced detailed semantic interpretations, it would have to know how to reason about partially restricted movement. Although this is a reasonable strategy, it is much simpler to coerce the notion of a limit into that of a barrier.

It would be possible to do away with the *ACCELERATE* and *DECELERATE* actions since their semantics mirror those of *ENABLE* and *BLOCK* respectively. They are included both to allow for the possibility of graded distinctions (although this exposition does not exploit this potentiality) and to make the coercion process more intuitive.

3.2 The Role of Syntax

For the examples of the previous subsection, the syntax of the sentence determines (i) what the path action-bearing constituents are and (ii) the order in which the path actions are to be applied to one another. Although the lexicon entry for a lexical item with a directional component must indicate its corresponding path action, (ii) can usually be determined from the syntactic category of the constituent that the lexical item ends up being a part of. For example, the direct object of a transitive verb is usually the target of the verb's path action. If the lexical item invokes a non-standard behavior or behaves differently in different syntactic situations, this tendency is indicated in its lexicon entry.

As the sentence is parsed, feature-structure representations [Shieber, 1986] are built. The constituent that immediately contains a lexical item with a path action is assigned the following features:

Path-action. This can take on one or more of the values *NEUTRAL, BLOCK, REMOVE-BLOCK, ENABLE, ACCELERATE, DECELERATE*. This feature indicates the directional contribution the encompassed lexical item makes to the interpretation of the sentence, and can vary depending on the syntactic category of the lexical item. For example, "lift" as a transitive verb can be a *REMOVE-BLOCK* type whereas as an intransitive verb it is *NEUTRAL*.

Target-entity. This indicates which constituent to apply the action to. It may be another action-bearing constituent, as is "the ban" in the phrase "to lift the ban on wastewater dumping" or it may be a constituent that is left unanalyzed, as "wastewater dumping" would be. A null value for this feature signals the end of an action application sequence.

Next-constituent. Usually this is unified to the value of the target-entity. However, there are cases when a constituent has more than one relevant argument, as in "The president shielded the elephants from attacks" where both "the elephants" and "from attacks" are complements of "shielded." In this case the value of the target-entity is retained until the end of the action application sequence and the current path-action is applied to the constituent that unifies with the next-constituent feature.

Viewpoint-agent. This is the entity, usually animate (metonymic agents such as "the White House," "Beijing," etc., are permissible here) and usually the subject of the main clause, whose opinion is being investigated. This feature is optional.

Secondary-agent. This attribute indicates a secondary agent that plays a role in the indication of the sentence's directionality. For example, in "The governor persuaded farmers to implement irrigation measures" it may be desirable to retain the information about who is doing the implementation, although often this is unstated. This feature is optional.

For some lexical items, the path-action assumed depends on the action found in the target-entity. For example, the verb *shield* acts as a *REMOVE-BLOCK* path-action in "The president shielded the elephants from attacks by poachers" because its next-constituent feature points to a constituent containing a *BLOCK* action. However, if the sentence were simply "The president shielded the elephants," the verb would act as an *ENABLE* action since both its target-entity and its next-constituent point to a constituent with a *NEUTRAL* action. Most lexical items that can take on the *REMOVE-BLOCK* action take on a different action when no *BLOCK*-type target-entity follows.

Most (open-class) lexical items will fall into one of a small set of categories, for example, *finite-transitive-remove-block-verb* or *block-noun*, so once an initial set is defined, classification of new words should be relatively simple. Closed-class items such as prepositions, can require special attention.

An important syntactic consideration that arises when determining the agent's attitude involves clausal attachment. Consider the following sentences:

(3a) The congresspersons introduced the bill that appeased the protestors.

(3b) The congresspersons voted against the bill that appeased the protestors.

(3c) The congresspersons introduced the bill (in order) to appease the protestors.

(3d) The congresspersons voted against the bill (in order) to appease the protestors.

(3e) The congresspersons voted against the bill to ban wastewater dumping.

In sentences (3a) and (3b) the restrictive clause modifies "the bill." Therefore, since the agent enables the bill in (3a), the agent favors the bill and also what the bill does. Similarly, in (3b) since the agent opposes the bill, the agent opposes what the bill does. Now consider (3c-d). In these sentences the infinitival clause acts as a purpose clause attached to the main verb, instead of a modifier for the object of the verb; this can have an affect on the interpretation. In these sentences the opinion of the congresspersons with respect to the protestors is independent of whether the congresspersons *ENABLE*d or *BLOCK*ed the bill. In both cases the goal is to appease the protestors. Thus when a purpose clause is identified, its contents should be preceded with an implicit *ENABLE* action.

Notice that in both (3a) and (3c) the main clause produces an *ENABLE* action and the attitude toward the complement is the same whether it is an object modifier or a purpose clause. However, when the main clause produces a *BLOCK* action, as in (3b) and (3d), the attitude to the contents of the complement can be affected. However, it is not always easy to distinguish the two cases, as shown by a comparison between (3d) and (3e), a variation of (1a). Although the sentences are syntactically quite similar, the first can be seen as having a purpose clause[1] and the second an object modifier. Unfortunately, determining the proper attachment requires more detailed interpretation than this model assumes. This issue will be discussed further in Section 5.

3.3 A Full Example

The following sentence is more complicated:

(4a) President Bush halted hardware sales

[1] Actually, sentence (3d) is ambiguous if "in order" is omitted, and if context is lacking.

(4b) to increase pressure on Beijing

(4c) after its crackdown on the pro-democracy movement.

The first conclusion is that President Bush is opposed to hardware sales because *BLOCK* $\triangleright$ *NEUTRAL* $\Rightarrow$ *BLOCK* $\Rightarrow$ oppose("Bush", "hardware sales") is easily induced from (4a). The next step is to recognize that (4b) is a purpose clause, thus indicating that again "President Bush" is the agent and that the sequence of path actions begins with an implicit *ENABLE*. The verb introduces an *ACCELERATE* action and its object is a *BLOCK* yielding *ENABLE* $\triangleright$ *ACCELERATE* $\triangleright$ *BLOCK* $\Rightarrow$ *ENABLE* $\triangleright$ *BLOCK* $\Rightarrow$ *BLOCK* $\Rightarrow$ oppose("Bush", "Beijing"). Thus President Bush opposes Beijing, or to interpret the metonymy, opposes some action of the government situated in Beijing. It would be useful to recognize the link between parts (4a) and (4b); namely that the hardware sales that Bush halted were sales to Beijing, but this is beyond the capability of the method as currently formulated.

Notice that in (4b) the noun "pressure" is followed by the preposition "on" which reinforces the indication that "Beijing" is the object of the *BLOCK* action. Prepositions in noun phrases are often strong directionality indicators. The prepositions "on" and "against" are associated with the *BLOCK* action, "from" with *REMOVE-BLOCK*, and "for" with *ENABLE*. The verb "protest" is interesting in this respect. Both "protest against X" and "protest about X" indicate a negative attitude, as does "protest X" with no preposition at all. However, "protest for" flips the polarity. This is another example of a syntactic consideration that can affect a lexical item's directionality.

The adjunct (4c) imposes some difficulties, the foremost being the determination of the referent of the pronoun. In general pronoun resolution requires a more sophisticated interpretation element than is assumed for this method (however, see Section 5). Assuming that this difficulty can be overcome, the matter of determining the agent's attitude toward the contents of the adjunct remains. While (4b) describes the desired results (increasing pressure) of the action of (4a) (halting sales), (4c) describes the justification or motivation for (4a) and (4b). In fact, (4c) can be paraphrased as "because Beijing cracked down on the pro-democracy movement."

Recall in sentences (3c-d), that the agent of the main clause favors the contents of the purpose clause. However, the agent of the paraphrased (4c) is one that has been opposed to in (4b). To handle this kind of situation, we need a heuristic that states that if in some sentence S agent A is found to oppose an entity E that happens also to be an agent, if E's actions are described in a purpose clause of S, then A opposes the actions described in this purpose clause. In order to accommodate this heuristic, the *BLOCK* action is prepended to the interpretation of (4c).

Although (4c) is a noun phrase, its analysis proceeds similarly to the verb phrases seen so far. In noun phrase interpretation, a modifier can affect the interpretation of its head noun. For example, in "further slaughter of elephants," "further" is associated with *ACCELERATE*, "slaughter" with *BLOCK*, and "elephant" with NEUTRAL, giving *ACCELERATE* $\triangleright$ *BLOCK* $\triangleright$ *NEUTRAL* $\Rightarrow$ *ACCELERATE* $\triangleright$ *BLOCK* $\Rightarrow$ *BLOCK* $\Rightarrow$ oppose(A, "elephants"). In the case of (4c) it is straightforward to associate *BLOCK* with "crackdown," but a judgement must be made as to whether to break "pro-democracy" into two pieces. Trying it the first way we get (after prepending a *BLOCK* for the pronoun, see above):

BLOCK $\triangleright$ *BLOCK* $\triangleright$ *ENABLE* $\triangleright$ *NEUTRAL* $\Rightarrow$
BLOCK $\triangleright$ *BLOCK* $\triangleright$ *ENABLE* $\Rightarrow$
BLOCK $\triangleright$ *BLOCK* $\Rightarrow$
ENABLE $\Rightarrow$
favor("Bush", "democracy movement")

If we hadn't broken "pro-democracy" in two, we would have ended up with favor("Bush", "pro-democracy movement").[2]

[2]In other situations, "movement" would be assigned an *ENABLE* path-action, but terms occupying a position at the end of the inference chain are generally left unanalyzed, in the interests of having something interesting to report. In other words, favor("Bush", E) is not very informative.

Discussion. The description in this subsection is deliberately high-level since the implementation and grammar (written in Common Lisp, based on a unification-based parser described in [Batali, 1991]) is an experimental prototype and covers only a few sentences. Furthermore, this exposition has omitted discussion of several important issues: how to handle negation, assignment of path actions to lexical items that occur in compounds, noun phrase disambiguation, interpretations of conjunctions, and so on. I leave all of this to future work.

3.4 The Role of General Metaphor in Path Action Assignment

How are lexical items assigned path actions? The approach outlined here is motivated by [Lakoff and Johnson, 1980], who observe that the use of "general metaphor" is structurally consistent and surprisingly widespread in "everyday" utterances. An example of a general English metaphor is one in which negative, undesirable things are described in terms of "downness" whereas desirable things are expressed in terms of "upness." This is evident in phrases such as "stocks took a dip," "the quality is declining," "it's going downhill," and so on. The central meaning of metaphors such as these can be considered domain independent, as evidenced by the fact that they are used in many diverse contexts. These observations about general metaphor are useful for deciding what path action to assign to a particular lexical item.

Words that are usually thought of as synonyms do not necessarily work the same way within these metaphoric structures. "Lift" is used to indicate the removal of some obstacle, whereas "raise" indicates an increase of a quantity. The phrase "raise the ban" is confusing to most readers and can even seem to mean the opposite of what "lift the ban" means. An appeal to general metaphor also helps explain the difference between sentences (1a) and (1b). Lakoff and Johnson identify what they call the Arguments Are Buildings metaphor. It is common to speak of an argument's foundations, whether or not a statement supports the argument or position, arguments that fall or collapse, constructing good arguments, and so on. This observation helps explain why "support" is classified as an *ACCELERATE* action.

The mechanism described here for assigning the value for the direction-type attribute is inspired by the idea of the pervasiveness of metaphoric extension. However, it does not strictly follow the structures that have been observed. Some metaphoric models are too specialized for the path metaphor, and some words participate in more than one metaphor and a choice among them, relevant to the path model, must be made.

4 Related Work

Related work in information retrieval and text categorization systems is discussed elsewhere in this volume. Therefore, this section focuses only on work related to the restricted interpretation model.

Like direction-based text interpretation, work in text skimming such as [DeJong, 1979, Lebowitz, 1983] involves extracting information from only certain parts of a text. These approaches differ from the method described here in that they rely heavily on domain-specific world knowledge. Furthermore, since they use little syntactic information they would have trouble making the directionality distinctions in, say, sentences (3b) and (3d). Jacobs ([Jacobs, 1990]) describes an approach called relation-driven text skimming that is similar in some ways to DTI. In both approaches, the relevant topics must be specified in advance and then only limited semantic interpretation is done and only on a subset of the sentences, chosen by a coarse first-pass. The main difference is that relation-driven skimming looks for the kind of information best indicated by predicate-argument relations (e.g., who is the target and who is the suitor of a corporate takeover). In order to determine these kinds of relations it makes more use of syntactic information than the other skimming methods, but the relation determination in many cases only requires a partial parse. This is advantageous since the parser need be less complicated than that needed by DTI, mainly because the kind of relation extracted is less subtly expressed.

The Plot Units strategy for text summarization [Lehnert, 1982] makes use of a distinction between positive and negative events. However, this distinction is made at the conceptual level (e.g., the fact that one's car won't start is a negative event) rather than at the level of a component in the interpretation of the meaning of a sentence. The polarity of these affect states are used to characterize a sequence of events in terms of a narrative primitive. For example, the sequence of a negative events motivating an action that terminates the cause of the negative event is a common sequence termed "resolution of a problem by intentional means." There may be a similar underlying motivation between this work and DTI but the actual mechanisms are quite disparate.

Several researchers (e.g.,[Carbonell, 1982] [Martin, 1990]) have integrated general, or conventional, metaphor into the process of full text understanding. The goals of these systems are different from that of DTI in that they use an understanding of the workings of metaphor in order to determine the author's intended meaning, instead of trying to coerce the meaning of a sentence into one all-purpose metaphor.

5 A Hybrid Text-Based Intelligent System

Direction-based text interpretation is meant to be a component of a hybrid TBIS system, although it is not currently implemented as such. In this section I attempt to flesh out the architecture that such a system might have.

First, the desired topical information (e.g., "environmental issues" or "wastewater dumping") is specified to an information retrieval system such as that described in Croft & Turtle (this volume). The parameters of the system are set such that sentence-level documents are returned (along with a few sentences of surrounding context) as candidates for directional analysis.

Next a robust partial parser like that of McDonald (this volume) processes the candidate sentences. The grammar is modified to use a feature-structure representation and the lexicon is augmented with the appropriate path action information. The conceptual analyzer is replaced by a module that interprets the resulting sentence structures in terms of the path model, as described in Section 3.

Since McDonald's parser is designed to be robust but partial, it may not produce some of the constituent attachment and pronoun resolution information that DTI is sensitive to. To remedy this problem while still avoiding the need for complex knowledge bases, we might try integrating results from statistical methods of text processing (see Stanfill, this volume). For example, [Hindle and Rooth, 1991] describes an approach for determing prepositional phrase attachments based on statistical tendencies in a large corpus. [Dagan and Itai, 1990] has done similar work on pronoun resolution. While this kind of approach to text analysis is still preliminary, it is also quite promising, and it seems likely that in future hybrid text interpretation systems will incorporate statistical results extensively.

The results of the analysis are returned to the user as predicate/document pairs, where the predicate has the form *direction(agent, event)*, e.g., favor("congresspersons", "wastewater dumping"). In addition to determining the answers to directional queries, the analysis component can act as a filter on the output of an information retrieval component in that if it cannot find directional content in the candidate sentence, then this information (and perhaps results from other kinds of analyses) should lower the relevance ranking of the document that contains it.

6 Conclusions

Ideally a Text-Based Intelligent System would perform full interpretation of its document corpus and allow the results to be accessed according to a user's information need. However, since the state of the art is quite far from this goal, much of the work in this volume suggests intermediate steps toward the ideal. These intermediate steps include: partial parsing (McDonald), partial representation (Hirst, Lewis), combination of weak interpretation methods (Wilks et al.), and statistical approximations to full understanding (Stanfill & Waltz, Maarek). This chapter is no exception: it proposes a question-answering paradigm that yields only a partial interpretation of a sentence's meaning.

More specifically, I have described a method for answering a general class of queries without engaging the complexity required by natural language processing techniques that attempt to generate "all plausible" inferences. This kind of approach is profitable only if the effort involved in building and executing the method does not outweigh the depth and quality of the results. If the effort does get too large, one could argue that a general text understanding system would be more appropriate since it could produce more detailed interpretations for the same amount of effort. In analyzing the tradeoffs of an approach like direction-based text interpretation, several questions need to be answered:

How often does the text contain the answer to the query in a form discernible to the method? In what situations must solutions to problems of anaphora, attachment, and lexical ambiguity must be resolved? How valid is the assumption that the target queries are general and useful enough to justify the effort required to answer them? And what kinds of direction-based queries, aside from the agent-attitude one explored here, can be answered using the proposed method?

Any restricted interpretation model must address these issues. The outcome of the tradeoffs can only be determined through empirical studies. If this approach and others like it can tip the balance in their favor, restricted semantic analysis will prove a useful component in the construction of efficient, intelligent text interpretation systems.

Acknowledgements. The majority of this research was completed during an internship at Xerox Palo Alto Research Center. The following people provided much useful discussion in the preparation of this chapter: Per-Kristian Halvorsen, John Batali, Susan Newman, Deborah Tater, Cathy Marshall, and Doug Cutting; additionally Robert Wilensky, Narciso Jaramillo, and Yoëlle Maarek provided helpful comments on various drafts. Finally, I would like to thank Paul Jacobs for advice on this draft and for continuous, generous encouragement.

References

[Batali, 1991] Batali, J. (1991). *Automatic Acquisition and Use of Some of the Knowledge in Physics Texts*. PhD thesis, Massachusetts Institute of Technology, Artificial Intelligence Laboratory.

[Carbonell, 1982] Carbonell, J. G. (1982). Metaphor: An inescapable phenomenon in natural-language comprehension. In Lehnert, W. G. and Ringle, M. H., editors, *Strategies for Natural Language Processing*, chapter 15, pages 415–434. Lawrence Erlbaum Associates, Publishers, Hillsdale, New Jersey.

[Dagan and Itai, 1990] Dagan, I. and Itai, A. (1990). A statistical filter for resolving pronoun references. In *Proceedings of the 7th Israeli Symposium on Artificial Intelligence and Computer Vision*.

[DeJong, 1979] DeJong, G. F. (1979). Skimming stories in real time: An experiment in integrated understanding. Technical Report Computer Science Department Technical Report #116, Yale University.

[Hindle and Rooth, 1991] Hindle, D. and Rooth, M. (1991). Structural ambiguity and lexical relations. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*.

[Jacobs, 1990] Jacobs, P. S. (1990). To parse or not to parse: Relation-driven text skimming. In *Proceedings of the 13th International Conference on Computational Linguistics*, volume 2, pages 194–198.

[Lakoff and Johnson, 1980] Lakoff, G. and Johnson, M. (1980). *Metaphors We Live By*. The University of Chicago Press.

[Lebowitz, 1983] Lebowitz, M. (1983). Memory-based parsing. *Artificial Intelligence*, 21:363–404.

[Lehnert, 1982] Lehnert, W. (1982). Plot units: A narrative summarization strategy. In Lehnert, W. G. and Ringle, M. H., editors, *Strategies for Natural Language Processing*, chapter 15, pages 375–414. Lawrence Erlbaum Associates, Publishers, Hillsdale, New Jersey.

[Lehnert and Sundheim, 1991] Lehnert, W. and Sundheim, B. (1991). A performance evaluation of text-analysis technologies. *AI Magazine*, 12(3):81–94.

[Martin, 1990] Martin, J. H. (1990). *A Computational Model of Metaphor Interpretation*. Academic Press, Boston.

[McCune et al., 1985] McCune, B., Tong, R., Dean, J., and Shapiro, D. (1985). Rubric: A system for rule-based information retrieval. *IEEE Transactions on Software Engineering*, 11(9).

[Reddy, 1979] Reddy, M. (1979). The conduit metaphor – a case of frame conflict in our language about language. In Ortony, A., editor, *Metaphor and Thought*, pages 284–324. University Press, Cambridge, England.

[Salton, 1988] Salton, G. (1988). *Automatic text processing : the transformation, analysis, and retrieval of information by computer*. Addison-Wesley, Reading, MA.

[Shieber, 1986] Shieber, S. M. (1986). *An Introduction to Unification-Based Approaches to Grammar*. C.S.L.I. Lecture Notes, Number 4. Center for the Study of Language and Information, Stanford University.

[Talmy, 1985] Talmy, L. (1985). Force dynamics in language and thought. In *Parasession on Causatives and Agentivity*, University of Chicago. Chicago Linguistic Society (21st Regional Meeting).