From: AAAI Technical Report FS-94-04. Compilation copyright © 1994, AAAI (www.aaai.org). All rights reserved.

Talking about Time: Temporal Reasoning as A Problem of Natural Language

Łucja Iwańska
Wayne State University
Department of Computer Science
Detroit, Michigan 48202, USA
lucja@cs.wayne.edu

Abstract

A new, natural-langauge style temporal reasoner capable of representing and reasoning with many previously unaccounted aspects of information conveyed by English temporal expressions, including information from arbitrary Boolean temporal expressions involving explicit *and*, *or*, *not* at different syntactic levels, temporal quantifiers and infinite number of temporal relations, is presented.

The paper demonstrates that important inferences about time can be captured by a general reasoning mechanism inherent to natural language, many aspects of which are mimicked by the proposed model of natural language. Interaction of language-inherent and time-specific computation is discussed.

1 Introduction

1.1 Temporal Reasoning– Critical and Widely Investigated Problem

Virtually all real-life problems involve handling temporal information, and temporal reasoning remains one of the most important and most widely investigated problems in artificial intelligence. In this paper, we show that reasoning about time is a problem in which natural language general reasoning capabilities mix in a very interesting way with the capabilities specific to understanding time. We demonstrate that time can be regarded as an object whose properties can be expressed and reasoned about via natural language in many ways analogous to reasoning about other objects such as books.

Within the field of artificial intelligence, temporal reasoning has not been considered as a problem of natural language. Originally, we planned to extend our Boolean algebra framework for natural language processing, the UNO model of natural language [Iwańska, 1994a] [Iwańska, 1994b] [Iwańska, 1993] [Iwańska, 1992a] [Iwańska, 1992b], by incorporating one of the very many existing temporal reasoners. In order to conduct their in-depth comparison and evaluate them with respect to the phenomena observed in natural language, we collected hundreds real-life examples of English temporal expressions of different syntactic categories[1]. We have closely examined these examples with respect to their information content and algorithms to compare this content, and concluded that temporal reasoning can be viewed as a problem of natural language because many of its important aspects can be accomplished by a general, natural language inherent inferencing mechanism. We also discovered that important aspects of temporal reasoning remain largely unadressed, and that our Boolean algebra framework provides answers to some of these problems.

1.2 Temporal Reasoner Natural Language Style

We have developed a new temporal reasoner which allows one to represent and reason with temporal information expressed in natural language. It automatically extracts explicit temporal expressions from textual documents and creates their semantic representation. This representation allows one to compute the entailed logical inferences and facilitates computing context-dependent implicatures, inductive-like inference common in natural language. For any set of temporal expressions, their information content can be computed and compared by computing semantic relation of entailment between their representations. This entailment relation can be used to compute answers to the *Yes-No* questions about time as well as answers to the *When ? How long ? How often ?* queries of the resulting knowledge base, and to compute temporal ordering of the events described in the documents. This entailment relation can also be used to automatically determine when to invoke a nondeductive reasoner that would make up for the system's knowledge gaps.

1.3 Novel Characteristics

Our temporal reasoner is novel in the following important ways:

1. It demonstrates that important inferences about time can be captured by a general reasoning mechanism inherent to natural language, many aspects of

[1]The sources of these data are several hundreds on-line *Wall Street Journal* documents, a few recent issues of the *Newsweek* magazine, a couple of short novels, excerpts from various books, and data found in various reference materials on English grammar.

which are mimicked by the UNO model of natural language.

2. It allows one to automatically represent and compute relative informativeness of English temporal expressions describing various aspects of time. This capability allows the system to automatically update its knowledge base as new information conveyed via natural language becomes available, to avoid storing redundant information, and efficiently handle temporal queries.

   While we share the goal of handling various aspects of temporal information conveyed by English expressions with [Kahn and Gorry, 1977] [Allen, 1984] [Allen, 1985] we propose a computational model whose formal and computational properties more fully reflect characteristics natural language, we consider a much wider spectrum of English temporal expressions, we systematically account for the differences in their meaning and inferences they license in the context of the current knowledge base via the representation computed automatically by the translation procedure motivated by the semantics and pragmatics of natural language, and finally, We attempt to account for reasoning about time with a general natural language inherent reasoning capabilities and identify those aspects of temporal reasoning that are outside of natural language.

3. It accounts for understanding and reasoning with information from arbitrary Boolean temporal expressions involving explicit *and*, *or*, *not* at different syntactic levels such as

   Example 1: *Not very often*

   Example 2: *Neither on Monday, April 22, 1992, nor five or six days ago*

   We are not aware of a temporal reasoner capable of handling general case of explicit negative information. Representing and reasoning with explicit negation in a semantically clean fashion is a unique characteristics of the UNO model [Iwańska, 1992a] [Iwańska, 1992b] [Iwańska, 1994a]. Limited forms of negation have been addressed by a number of researchers, including the Allen's landmark algebra of time intervals [Allen, 1984] [Allen, 1985]. Temporal reasoners often inherit problems with negation from the assumed representation. For example, [Dean and McDermott, 1987] implemented their temporal reasoner in PROLOG, which allows them to represent propositional (sentential) negation only, and offers lacking semantics negation as failure. A number of temporal reasoners are capable of handling disjunctive information, usually its various limited forms [Allen, 1985] [Shoham, 1987] [Schrag et al., 1992] [Gerevini and Schubert, 1993].

   We believe that the UNO capability of expressing and propagating arbitrary Boolean constraints can significantly reduce an average computational cost of updating temporal aspects of a knowledge base and handling temporal queries because such expressions allow one to fully state known information; disjunction often expresses possible choices or outcomes, and negation often reflects ruling them out.

4. It handles temporal quantifiers such as

   Example 3: *Sometimes, but not always*

   Example 4: *Not very often*

5. It handles many more temporal relations observed in natural language than any other temporal reasoner, for example,

   Example 5: *Some time after*

   Example 6: *Five or six days before*

   Example 7: *Not long ago*

   It also gives a different meaning to the *overlap* relation that allows one to capture incomplete knowledge with respect to the extent of the overlap (much like in the set theory) such as *overlaps partially only*, *contained in*, and *equal to*.

6. It allows one to reason with both qualitative and quantitative temporal information. For example, it can compute logical entailments of the sentences

   Example 8: *Mark is 25*

   Example 9: *Mark is not old*

1.4 Current Research

Our current efforts concentrate on three major aspects of temporal reasoning:

1. Identifying and handling context-dependent, nondeductive aspects of temporal reasoning.
2. Expressing qualitative and quantitative temporal information one with another.
3. Interplay of the semantics and pragmatics of temporal inferences.

1.5 Long Term Research Goal: Most Adequate Computational Model of Natural Language

The presented temporal reasoner is part of our effort to create the most adequate computational model of natural language that would capture all aspects of information conveyed via natural language, and closely simulate its representation scheme and inferencing mechanism. The significance of this line of research stems from the fact that today myriads of on-line textual documents are available, yet only an extremely small portion of the information contained in them is utilized and put together. A successful model of natural language will allow programs to automatically create knowledge bases closely resembling knowledge a person would gain after reading the documents; these knowledge bases can be subsequently queried by the same computational mechanism closely resembling human reasoning in natural language that was responsible for their creation. Sharing all aspects of information, its representation, and inferencing scheme by the information source and the quering party assures little, if any, information loss.

1.6 This Paper

The emphasis of this paper is demonstrating that temporal reasoning can be viewed as a problem of natural language. We attempt to outline the entire problem of reasoning about time in natural language, discuss the connections with different areas of natural language processing and other areas of artificial intelligence. Due to space limitations, we include very few technical details (they can be found in [Iwańska, 1994c]), but present many examples that illustrate the discussed problems and unique capabilities of our natural-language-style temporal reasoner.

The rest of this paper is organized as follows: Section 2 discusses natural language inherent reasoning about time; Section 3 discusses time-specific knowledge and inference; Section 4 contains conclusion.

2 Natural Language Inherent Reasoning about Time

2.1 Logical, Context-Independent Inference

If one is told that

Example 10: *Long before the market crash*, *Sam lost a lot of money*

then one automatically understands that it is also the case that

Example 11: *Sam lost a lot of money before the market crash*

Understanding entailment of these sentences results from understanding that the expression *long before* contains more information than the expression *long*. Understanding such semantic relations involving temporals is exactly analogous to the human understanding of natural language expressions describing properties of various objects, including complex adjectives, common nouns, and verbs. Understanding and reasoning with such complex natural language references to time is exactly analogous to handling complex references to other objects. For example, the sentence

Example 12: *Sam is a short woman*

entails the sentence

Example 13: *Sam is a woman*

because the noun phrase *a short woman* contains more information than the noun phrase *a woman*; this is because a common noun *short woman* contains more information than the noun *woman*; the UNO representation of these noun phrases guarantees to preserve this relation. The same existing UNO representation and inference mechanism can be used to utilize information from and compute referents of such complex references to time as

Example 14: *It was an extremely difficult and unhappy time*

For example, the UNO system can infer from this sentence that it is also the case that

Example 15: *It was a difficult time*

Example 16: *It was not a very happy time*

Example 17: *It was not an easy time*

The most characteristic property of negation in natural language, that it reverses entailment, is also exhibited by temporals: the sentence

Example 18: *Sam is not a woman*

entails the sentence

Example 19: *Sam is not a short woman*

because *not a woman* entails *not a short woman*; exactly analogously, the sentence

Example 20: *Sam lost a lot of money, but not before the market crash*

entails the sentence

Example 21: *Sam lost a lot of money, but not long before the market crash*

because *not before* entails *not long before*.

Another example of natural language inherent reasoning capabilities involved in reasoning about time is understanding that the sentence

Example 22: *He never smiles*

logically entails the sentence

Example 23: *It is not very often that he smiles*

which is based on understanding that a quantification frequency adverb *never* logically entails *not very often*.

2.2 Scalar Nature of Temporals

Logical entailments such as between *never* and *not very often* are due to the scalar predicate nature of the quantification frequency adverbs. Scalar predicates, a large, cross-categorial class of expressions of natural language that denote certain quantitative values, also includes cardinal numbers, determiners, evaluative and gradable adjectives, modals, and epistemic verbs. For scalar predicates, semantic relation of entailment holds between concepts that are denotations of lexically simple scalar predicates. It is this entailment relation between *always* and *sometimes* that accounts for the fact that when told

Example 24: *John always smiles*

one automatically understands that it is also the case that

Example 25: *John sometimes smiles*

Qualitative Scales, one of the knowledge representation formalisms underlying the UNO model, captures semantic properties of scalar predicates and allows one to uniformly compute logical entailments between arbitrary scalar expressions of different syntactic categories. The representation of complex scalar expressions such as

Example 26: *sometimes, but not very often*

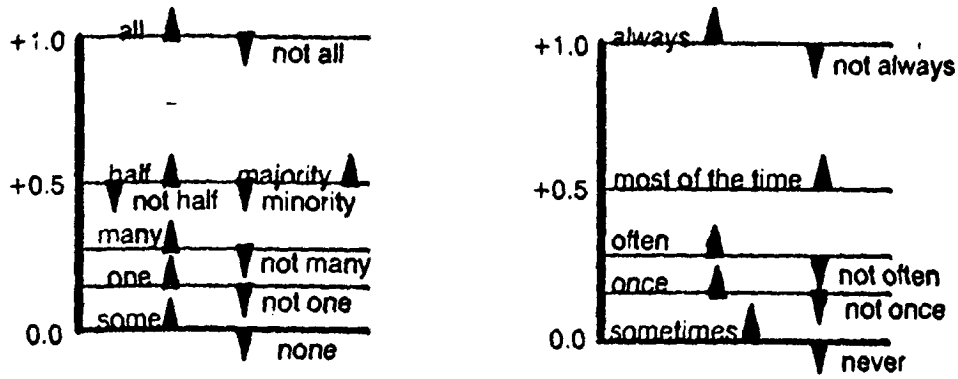

Fig. 1. Scalar semantic properties of determiners (the left scale) are paralled by frequency quantification adverbs (the right scale).

is the result of the Boolean operations of *meet*, *join*, and *complement* on the representation of lexically simple scalars. These semantically clean operations model scalar aspects of conjunction, disjunction, and negation in natural language, and allow one to compositionally derive semantic representation of complex Boolean expressions of scalar predicates. In the UNO model, computing entailment between complex adverbs

Example 27: *never*

Example 28: *not very often*

and between complex adjectives

Example 29: *terrible*

Example 30: *not very good*

is accomplished by the same mechanism.

Figure 1 gives two examples of qualitative scales:[2] determiners and frequency quantification adverbs. These scales are very similar and reflect the fact that with respect to some logical inference, temporal quantifiers behave exactly the same as determiners that give rise to the generalized quantifiers of natural language. The same distinction between first-order and non-first-order definable quantifiers applies: the adverb *always* behaves like the determiner *all*, and the adverb *often* like the determiner *many*.

Many other scalars are also applicable to describing time. Figure 2 gives four examples of such scales. The upper left scale of evaluative adjectives and the lower left scale of intensifying adverbs allow the UNO model to account for the entailment between the temporal noun phrases

Example 31: *a very good time*

Example 32: *not a bad time*

and between the temporal noun phrases

Example 33: *a terrible time*

[2]Only lexically simple scalars and their negations can be conveniently illustrated by such a picture.

Example 34: *not a very good time*

in the exact same fashion as for the entailment between the nontemporal noun phrases

Example 35: *a very good book*

Example 36: *not a bad book*

and between the nontemporal noun phrases

Example 37: *a terrible student*

Example 38: *not a very good student*

The upper right scale in Figure 2 relates words describing degree of similarity of some objects. This scale allows the UNO model to mimic human understanding of natural language and infer from the sentence

Example 39: *X and Y are* ***the same***

that

Example 40: *X and Y are* ***very similar***

Example 41: *X and Y are* ***not very different***

Understanding logical implications of sentences in which the degree of similarity words describe temporal information is exactly analogous. The UNO system infers from the sentence

Example 42: *X and Y happened* ***at the same time***

that it is the case that

Example 43: *X and Y did not happen* ***at a different time***

The last scale in Figure 2, in the lower right corner, allows the UNO system to infer from the sentence

Example 44: *It took* ***an extremely long time***

the falsity of the sentences

Example 45: *It took* ***a very short time***

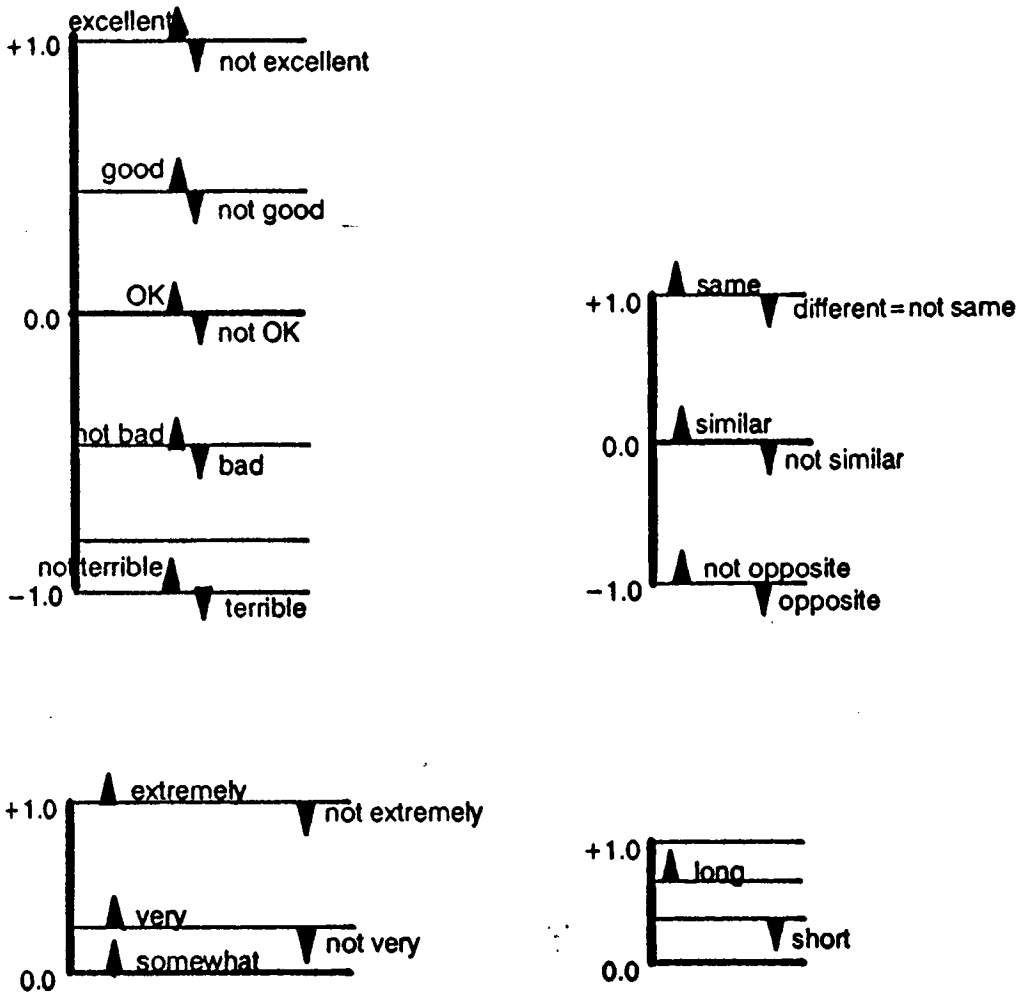

Figure 2: Scalars relevant for expressing temporal information

Example 46: *It did not take* [*a long time*]

The same UNO mechanism accounts for the relevant logical inferences when these adjectives describe nontemporal objects. The UNO system infers that if

Example 47: *Pluto has* [*an extremely short tail*]

then it is also the case that

Example 48: *Pluto does not have* [*a very long tail*]

Some scalar predicates convey temporal information only. Figure 3 gives examples of two such scales. They allow the UNO model to account for the entailments between the following pair of sentences:

Example 49: *John is* [*very old*]

Example 50: *John is* [*not young*]

and represent and reason with the information from exchanges like:

Example 51: *Was John* [*late*] *for the meeting* ?

Example 52: *No, this time, he showed up* [*very early*]

When inferencing with sentences involving scalar predicates, the formal model of Qualitative Scales offers two major advantages over the concept of linguistic variable [Zadeh, 1975a] [Zadeh, 1975b] [Zadeh, 1975c] (a detailed comparison can be found in [Iwańska, 1992b] [Iwańska, 1994b]):

First, in case of linguistic variable, in order to compute entailment between expressions such as *very old* and *not young*, these expressions have to be converted

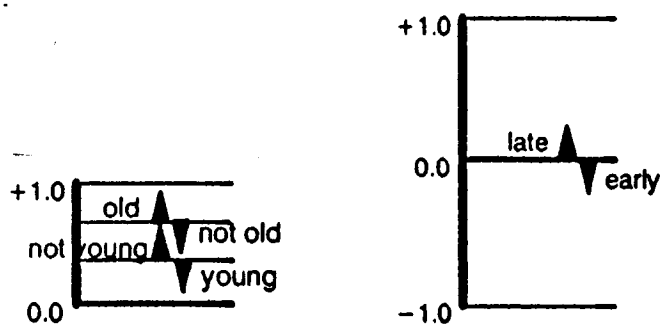

Figure 3: Scalars conveying time-specific information only

to numbers; this is because linguistic variable is not entirely qualitative, but rather it is about numbers; Qualitative Scales can be entirely qualitative, and inference from *very old* to *not young* can be accomplished without involving any numbers; this is important because underlying numeric values may not be available or may be unimportant; at the same time, Qualitative Scales can also be about numbers, should numeric values be available and desirable.

Second, meaning of the entailment relation in the UNO model reflects relative informativeness of natural language expressions– the entailed expression contains less or equal information than the expression by which it is entailed. This semantics closely corresponds to the semantics of natural language sentences with scalar expressions. For example, the scalar verb *like* contains less information than *love* and therefore the sentence

Example 53: *He* [*loves*] *cookies*

logically entails the sentence

Example 54: *He* [*likes*] *cookies*

In case of linguistic variable, the entailment relation between *excellent* and *good*, if any [3], reflects the fact that for each possible value of the relevant variable, one expression is more compatible with it than the other. This meaning of the entailment relation seems totally disjoint from the meaning of natural language sentences involving scalar expressions.

[3]In fuzzy logic, specific values are arbitrary– user can define these words to be partially or fully compatible, or incompatible.

2.3 Uncertainty at Different Syntactic Levels

We have recently expanded the UNO model to handle uncertainty at different syntactic and semantic levels in order to correctly preserve information present in natural language. Consider the following two sentences:

Example 55: *It is* [*possible, but not very likely,*] *that it was John who stole the bike*

Example 56: *John is* [*likely, but not extremely, likely*] *to be angry*

The first sentence expresses the uncertainty about the agent of the *bike stealing* action, and considers John as a *possible, but not very likely* candidate; semantically the likelyhood phrase *possible, but not very likely* modifies the subject noun phrase *John*.

The second sentence describes the uncertainty about John possessing the *angry* property, and considers this property *likely, but not extremely, likely*; semantically the likelyhood phrase *likely, but not extremely, likely* modifies the verb phrase *be angry*.

Time-related uncertainties can also be expressed at different syntactic levels, some of them quite complex as illustrated by the sentence

Example 57: *It is* [*very likely*] *that it happened shortly* [*before or after*] *the meeting*

In this sentence, the likelyhood phrase *very likely* modifies a sentence involving yet another uncertainty expressed by the disjunctive phrase *before or after*; the uncertainty expressed via explicit disjunction pertains the temporal relation between the time of the meeting and

```
temporal_relation ==>
   [ same_time,  different_time ] .

different_time ==>
   [ common_part,  no_common_part ] .

common_part ==>
   [ overlap ] .

overlap ==>
   [ begin, end, during ] .

no_common_part ==>
   [ after,  before,
     follow, precede,
     prior,
     later,  earlier,
     between,
     within
   ] .
```

```
after == follow .
after == later .

before == precede .
before == earlier .
before == prior .

same time == overlap fully .
same time == coincide with .
same time == simultaneous .
same time == concurrent .

between X and Y == after X and before Y
```

Figure 4: Hierarchy of basic temporal relations and their equivalences

the time of the event referred to by the second pronoun *it*.

2.4 Nonlogical, Context-Dependent Inference

Scalar temporals also engage in an extremely common in natural language pragmatic, context-dependent, inductive-like inference of scalar implicature [Horn, 1972] [Horn, 1989] [Hirschberg, 1985] [Reiter, 1990] [Iwańska, 1994b]). For example, the sentence

Example 58: *He* $\boxed{\textit{rarely}}$ *washes dishes*

may be an understatement nonlogically implicating that actually

Example 59: *He* $\boxed{\textit{never}}$ *washes dishes*

in the exact same fashion as the sentence

Example 60: *It is* $\boxed{\textit{not very good}}$

may be a polite way of saying that

Example 61: *It is* $\boxed{\textit{terrible}}$

The same interplay of the semantics and pragmatics of natural language, when a logically weaker sentence conveys the information of a logically stronger sentence, can be observed for temporal expressions. For example, the sentence

Example 62: $\boxed{\textit{Sometimes}}$ *he sings*

is logically consistent with

Example 63: *He* $\boxed{\textit{always}}$ *sings*

but may nonlogically implicate that

Example 64: *It is* $\boxed{\textit{not always}}$ *that he sings*

This way a logically weaker sentence *Sometimes he sings* may convey the information of the logically stronger sentence *Sometimes, but not always, he sings*

2.5 Other Complex References to Time

References to time can be even more complex and consist of many page description of certain events ending with a sentence like

Example 65: *Those were the happy times*

In such descriptions, time intervals are being characterized by the events that took place during those intervals, customs, trends in science, fashion, etc. Representing and reasoning with such complex references to time appears to require a full computational model of natural language. This provides an additional argument for our natural-language-as-an-extremely-expressive-formal-language line of research. The more complex subset of natural language our UNO model will be capable of representing, the more information about time the system will understand and utilize.

2.6 Absolute Time

Temporal information can be given in absolute terms via a description of certain time interval. Very common, particluar in factual reporting texts is a date, typically consisting of three elements: day, month, year. Due to different conventions about the presentation order, the same expression can refer to a different time interval; for example, in the USA

Example 66: *6.12.94*

refers to the 12-th day of the 6-th month (June) of the year 1994; in Europe the same expression refers to the 6-th day of the 12-th month (December) of the same year.

Time of some events can be provided in terms of qualitatively described intervals. For example, the sentence

Relation	Examples
same-time	*The meetings took place* [*at the same time*] *These* [*simultaneous*] *events surprised many* *The same basic machines will be sold* [*simultaneously*] *in the United States, Europe and Japan* *He named ambassador to New Zealand Takeo Iguchi to serve* [*concurrently*] *as ambassador to western Samoa* *The announcement of the pact* [*coincided*] *with the release by Volkswagen of its second-quarter earnings report*
start	*His talk* [*starts*] *the conference* *The conference* [*starts*] *with his talk*
end	*His talk* [*ends*] *the conference* *The conference* [*ends*] *with his talk*
during	*Operations came to a halt* [*during*] *the Cultural Revolution*
after	*Dancing was* [*after*] *singing* [*After*] *singing they danced*
before	*Singing was* [*before*] *dancing* [*Before*] *dancing, they sang*
follow	*Dancing* [*followed*] *singing* *First it was singing; then dancing* [*followed*]
precede	*Singing* [*preceded*] *dancing*
prior	*Singing was* [*prior*] *to dancing*
later	*First it was singing; dancing was* [*later*]
earlier	*Now they dance; singing was* [*earlier*]
between	*He was born sometimes* [*between*] *the First and Second World Wars*

Table I: Basic temporal relations: time of different events can be related by simple natural language expressions.

Example 67: *Ann took her exam* [*early in the day*]

describes the time of Ann's exam via the expression *early in the day* that refers to a particular, qualitatively described part of an time interval referred to by the expression *the day*.

Temporal intervals can also be described by events that happen during these intervals. For example, the sentence

Example 68: [*The early decades of this century*] *are characterized by dramatically encreased emmigration*

describes the interval referred to by the expression *The early decades of this century*.

2.7 Relative Time

Temporal information is often given in relative terms with respect to the time of some event. Important and relevant events often provide such a frame of reference. There are a number of simple natural language expressions that allow one to relate time intervals of different events.

Figure 4 shows hierarchy of such basic temporal relations arranged with respect to the extent of the overlap of the intervals. Many of these basic relations convey similar information. This similarity is encoded in the UNO system as type equalities also shown in this figure.

The number of possible temporal relations that can be expressed in natural language is infinite because basic relations can be further specified via complex expressions involving modifiers, Boolean connectives *and*, *or*, *not*, and cardinal numbers. Tables 2.2 and 2.2 give examples of natural language sentences involving basic and complex temporal relations.

Some of the relations can also be used to describe a temporal relation between some event and some temporal interval. For example, the sentence

Example 69: *Paul gave a talk* [*earlier this summer*]

places Paul's talk at a temporal interval denoted by the expression *earlier this summer*.

Temporal intervals can also be described in terms of their boundaries, as illustrated by the sentence

Example 70: *It didn't start* [*until 5 pm*]

Relation	Modifier	Examples
after	five months or so	*Five months or so after* | *the treatment, he was fully recovered*
	few	*It happened* | *a few minutes after* | *their conversation*
	not, very long	*He fell asleep* | *not very long after* | *the meeting*
follow	immediately	*The arrest* | *immediately followed* | *his confession*
before	five or six days	*Five or six days before* | *the trial, he confessed*
after, before	or, shortly	*He gave me these papers* | *shortly before or after* | *the meeting*
after	not, immediately	*It* | *didn't* | *happen* | *immediately after* | *the exam*
earlier	not, much	*John came home* | *not much earlier* | *than Mark*
	two years	*Two years later*, *she was married*
	this month	*The announcement was made* | *earlier this month*
follow	shortly	*Dancing* | *followed shortly*
during	not	*It didn't happen* | *during* | *the party*

Table II: Complex temporal relations: time of different events can be related by complex expressions of natural language; these complex relations compositionally derive from basic temporal relations via modifiers, Boolean connectives *and*, *or*, *not*, and cardinal numbers.

2.8 Temporal Anaphora

Some useful inferences can be drawn without resolving temporal definite anaphora. For example, when presented with the sentence

Example 71: *It happened* | *early in the day* |

it is not necessary to understand which day is *the day* in order to correctly answer the question

Example 72: *Did it happen* | *late in the day* | *?*

because the answer to this question can be inferred from the semantic relation between *early* and *late*.

Handling other questions may crucially involve resolving temporal anaphora in order to gain the necessary information. It appears that many types of temporal anaphora can be resolved with slightly modified fairly standard definite anaphora resolution algorithms. However, the general problem of computing definite references to time remains open because it crucially depends on the capability of truthfully representing all aspects of information contained in natural language texts.

One problem with resolving definite anaphora, including temporal anaphora, is its prohibitive computational cost. In our temporal reasoner, we have adopted the strategy of fully computing referents of definite anaphora only if some question of the knowledge base cannot be handled without it.
contains a definite anaphora *then* whose referent can only be resolved in descriptive terms.

Example 73: *since the early decades of this century*

2.9 Temporal Idioms

In the UNO model, information content of temporal idioms is approximated with information contained in expression s whose meaning and representation is dervied compositionally. For example, the UNOP system takes the meaning of the expression

Example 74: *now and then*

to be equivalent to the meaning of the complex expression

Example 75: *sometimes, but not very often*

2.10 Qualitative and Quantitative Temporal Information

The UNO model can represent and reason with both qualitative and quantitative temporal information such as

Example 76: *John is* | *very old* |

Example 77: *John is* | *25* |

What it cannot do is to automatically express one with the other. One difficulty stems from the fact that such mutual conversions necessarily involve handling some intensional aspects of natural language. For example, if John is 25, he may be a young researcher, but an old skaut.

The capability of automatically expressing qualitative and quantitative information one with another would be very useful in improving cooperativeness of a knowledge base. For example, if told that

Example 78: *X occurred* | *on May 16, 1991* |

a cooperative knowledge base is expected to answer a question like

Example 79: *Did X happen last year ?*

with something like

Example 80: *It happened on May 16, 1991*

a behaviour that the UNO model currently exhibits. However, if the knowledge base contained a more specific information such as

Example 81: *X occurred five days before their first press conference, on May 16, 1991, at 5:06:12 pm*

then for some users an answer with all these details may be irrelevant and very annoying because in the context of their knowledge and goals a cooperative answer that a user might have expected would be something like

Example 82: *It happened not long before Y*

where *Y* refers to an event important for them.

The qualitative and quantitative temporal expressions in the two above sentences contain information that only partially overlaps, and inference necessary to replace the first expression with the second expression is of both deductive and inductive nature. Inference necessary for achieving this kind of cooperation is deductive because a differently packaged information cannot contradict the knowledge base. This inference is inductive because the conditions of replacing the absolute expression with a relative expression depends on context and different notions of how long is long.

3 Special Knowledge about Time

When reasoning about time, people also employ knowledge necessary to understand temporal information that seems to be more like temporal domain knowledge and reasoning than natural language inherent capabilities.

3.1 Temporal Units

One example is knowledge about different units of time such as *day* or *month*, their semantic relations, eg. that the unit of *month* is larger than the unit of *day*, and quantitative relations such as the fact that a given month consists of a particular number of days. Such quantitative and qualitative relations between different units of time are encoded in the UNO model by the type equations of the following form

```
1 Larger_Unit == Number Smaller_Unit .
```

For example,

```
1 century == 100 year .
1 century == 10  decade .
1 year    == 12  month .
```

Such equations are used for numerical conversions between different units, as well as for inferring qualitative ordering relation on temporal units, with both aspects contributing to computing entailments of sentences. For example, it can be computed that the sentence

Example 83: *It took him* *three days*

entails the sentence

Example 84: *It took him* *two or three hours*

because the denotation of *three days* entails the denotation of *two or three hours*.

Modifiers, numbers and Boolean connectives allow one to create complex units of time of arbitrary length as illustrated by the temporal expressions in the following sentences

Example 85: *It took me* *one year, two months, three weeks, and four and a half day.*

Example 86: *Large part of the day*, *he was busy grading the exam*

3.2 Context-Dependent Grain Size

.2ex[Golding et al., 1992] found that in some contexts, particular units of time may be most appropriate when answering questions about future events. In their experiments, human subjects in response to the question

Example 87: *When is the meeting on salary increases ?*

produced answers that included such units of time as *hour* and *minute*; the answer such as

Example 88: *49,000.456 miliseconds*

was never produced.

However, given incomplete knowledge and therefore wrong or no expectations, a temporal reasoner must be capable of understanding exchanges involving arbitrarily different grains of temporal units such as the following exchange in which the question

Example 89: *How many months does it take for an electron to travel through the channel ?*

answered by an ignorant in physics student with

Example 90: *Two months*

instead of expected

Example 91: *Two nanoseconds*

3.3 Intensionally Defined Temporal Units

Some words such as *instant* and *while* denote very small time units; others such as *eon* and *eternity* refer to very long intervals. Meaning of such intensional units is represented in the UNO model as type equations created automatically from their dictionary definitions. For example, on page 730 of the Second Edition of the Webster's New World Dictionary gives the following definitions

Instant– a point or very short space of time, moment

Eon– an extremely long, indefinite period of time

The UNO representation of such definitions truthfully preserves the information they contain, and allows the system to draw relevant inferences. For example, they allow the system to conclude that if some event took an instant, then this event did not take a long time.

3.4 Propagating Temporal Constraints

Propagating temporal constraints often amounts to computing some transitive closure, and also appears not to be about natural language, or at least, not entirely about natural language. Obviously, the entailments illustrated by the examples in the previous sections are part of computing (propagating) temporal constraints. But equally obviously, computing consistency of partial temporal orderings [Allen, 1985] [Davis and Carnes, 1991] [Dechter et al., 1991] takes computation outside natural language.

The UNO model of natural language with its general purpose knowledge representation formalisms allows one to represent and propagate such constraints in a similar fashion as it can be done in PROLOG [Sterling and Shapiro, 1986], and offers additional bonuses of handling explicit negative information and a variety of quantifiers.

We have currently adopted a convention of computing transitive closure-like consistencies only per explicit request, partially because it is so expensive, and partially because people do not seem to perform such consistency checking unless they are forced to.

3.5 Deductive versus Nondeductive Inference

Except for scalar implicature, the UNO temporal reasoner is a deductive system acting as a conservative reasoner that does not jump to inductive, unsupported conclusions. While it can reason with incomplete or incorrect information and derive some useful conclusions, it does not make up for the gaps in its knowledge. However, it is capable of automatically identifying such gaps; this happens, for example, when the answer to a query is not entailed by the knowledge base, or when externally provided information (from some user or from some textual document) contradicts its current knowledge.

Any practical system must be capable of creatively overcoming its incomplete knowledge and smartly jumping to inductive conclusions, so we have built our temporal reasoner in such a way that any nondeductive (nonmonotonic) system can be imposed on top of it. We currently investigate the problem of which nondeductive system is most suitable for handling reasoning about time in natural language.

3.6 Causality and Domain-Specific Temporal Relations

In case of causality and domain-specific temporal relations, we again benefit from our natural-language-as-an-extremely-expressive-formal-language approach. We supply this information in the form of textual descriptions whose UNO representation is automatically computed, added to the knowledge base, and reasoned with. We currently extend the English coverage of the UNO model in order to automatically represent knowledge reflecting causality and other temporal constraints, and be able to automatically understand these constraints from such complex sentences as

> **Example 92**: *It is not uncommon for our new foreign students to take one or two crush courses on English language during their first semester*

4 Conclusion

We have shown that reasoning about time is a problem in which natural language general reasoning capabilities mix in a very interesting way with the capabilities specific to understanding time. We have demonstrated that important inferences about time can be captured by a general reasoning mechanism inherent to natural language, many aspects of which are mimicked by the UNO model of natural language.

We have presented a new temporal reasoner capable of representing and reasoning with many previously unaccounted aspects of temporal information conveyed by English temporal expressions, including reasoning with information from arbitrary Boolean temporal expressions involving explicit *and*, *or*, *not* at different syntactic levels, handling temporal quantifiers and infinite number of temporal relations.

References

[Allen, 1984] Allen, J. F. (1984). Towards a General Theory of Action and Time. *Artificial Intelligence*, 23:123–155.

[Allen, 1985] Allen, J. F. (1985). Maintaining knowledge about temporal intervals. In *Readings in Natural Language Processing*, pages 509–522. Morgan Kaufmann Publishers, Inc.

[Davis and Carnes, 1991] Davis, W. and Carnes, J. (1991). Clustering Temporal Intervals To Generate Reference Hierarchies. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR91)*, pages 111–117.

[Dean and McDermott, 1987] Dean, T. L. and McDermott, D. V. (1987). Temporal Data Base Management. *Artificial Intelligence*, 32:1–55.

[Dechter et al., 1991] Dechter, R., Meiri, I., and Pearl, J. (1991). Temporal Constraint Networks. *Artificial Intelligence*, 49:61–96.

[Gerevini and Schubert, 1993] Gerevini, A. and Schubert, L. (1993). Efficient Temporal Reasoning through Timegraphs. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-93)*.

[Golding et al., 1992] Golding, J., Magliano, J., and Hemphill, D. (1992). WHEN: A Model for Answerinh "When" Questions About Future Events. In Lauer, T., Peacock, E., and Graesser, A., editors, *Questions and Information Systems*. Lawrence Erlbaum Associates.

[Hirschberg, 1985] Hirschberg, J. B. (1985). A theory of scalar implicature. PhD thesis, Technical Report MS-CIS-85-56, Dept. of Computer and Information Science, University of Pennsylvania.

[Horn, 1972] Horn, L. R. (1972). *On the Semantic Properties of Logical Operators in English*. IULC.

[Horn, 1989] Horn, L. R. (1989). *A Natural History of Negation*. The University of Chicago Press.

[Iwańska, 1992a] Iwańska, Ł. (1992a). A General Semantic Model of Negation in Natural Language: Representation and Inference. In *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (KR92)*, pages 357–368.

[Iwańska, 1992b] Iwańska, Ł. (1992b). *A General Semantic Model of Negation in Natural Language: Representation and Inference.* PhD thesis, Dept. of Computer Science, University of Illinois at Urbana-Champaign. Also available as Technical Report UIUCDCS-R-92-1775 or UILU-ENG-92-1755.

[Iwańska, 1993] Iwańska, Ł. (1993). Logical Reasoning in Natural Language: It Is All About Knowledge. *International Journal of Minds and Machines, Special Issue on Knowledge Representation for Natural Language*, 3(4):475–510.

[Iwańska, 1994a] Iwańska, Ł. (1994a). A Boolean Algebra Approach to Reasoning with Explicit Negative Information in Natural Language (57 page article currently under review). *Journal of Theoretical Computer Science.*

[Iwańska, 1994b] Iwańska, Ł. (1994b). Qualitative Scales: Modeling Boolean Semantics of Scalar Predicates (40 page article currently under review). *Journal of Computational Linguistics.*

[Iwańska, 1994c] Iwańska, Ł. (1994c). Reasoning about Time in Natural Language. *In preparation.*

[Kahn and Gorry, 1977] Kahn, K. and Gorry, G. A. (1977). Mechanizing Temporal Knowledge. *Artificial Intelligence*, 9:87–108.

[Reiter, 1990] Reiter, E. (1990). The computational complexity of avoiding conversational implicatures. In *28-th Annual Meeting of the Association of Computational Linguistics.*

[Schrag et al., 1992] Schrag, R., Boddy, M., and Carciofini, J. (1992). Managing Disjunction for Practical Temporal Reasoning. In *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (KR92)*, pages 36–46.

[Shoham, 1987] Shoham, Y. (1987). Temporal Logics in AI: Semantical and Ontological Considerations. *Artificial Intelligence*, 33:89–104.

[Sterling and Shapiro, 1986] Sterling, L. and Shapiro, E. (1986). *The Art of Prolog.* The MIT Press, Cambridge.

[Zadeh, 1975a] Zadeh, L. A. (1975a). The Concept of a Linguistic Variable and its Application to Approximate Reasoning– I. *Information Sciences*, 8:199–249.

[Zadeh, 1975b] Zadeh, L. A. (1975b). The Concept of a Linguistic Variable and its Application to Approximate Reasoning– II. *Information Sciences*, 8:301–357.

[Zadeh, 1975c] Zadeh, L. A. (1975c). The Concept of a Linguistic Variable and its Application to Approximate Reasoning– III. *Information Sciences*, 9:43–80.