

LEARNING TEMPORAL ANNOTATION OF FRENCH NEWS

A Thesis
submitted to the Faculty of the
Graduate School of Arts and Sciences
of Georgetown University
in partial fulfillment of the requirements for the
degree of
Master of Science
in Linguistics

By

Jennifer A. Baldwin, B.A., B.S.

Washington, DC
April 22, 2002

LEARNING TEMPORAL ANNOTATION OF FRENCH NEWS

Jennifer A. Baldwin, B.A., B.S.

Mentor: Inderjeet Mani, Ph.D.

ABSTRACT

Annotation of temporal information in text can be a complex task for both humans and computers. Impressive systems exist that tag and annotate temporal information within a single language, but in order to adapt them to a different language, they require a considerably expensive investment of time and human knowledge. In this paper, I present a language-independent temporal annotation system that learns from an annotated corpus rules that it can apply to raw data. Evaluation of the system on news text shows promising results in multiple languages, supporting the applicability of this approach for the purpose of building annotation systems for new languages at minimal time and cost.

TABLE OF CONTENTS

Introduction	1
Approach	4
Corpora	4
Methodology	6
Background	6
Learning Component	15
Limitations	22
System Architecture	23
Results	24
Rules	24
PAST_REF, PRESENT_REF, and FUTURE_REF	30
Frequencies	32
Durations	34
Full and Partial Dates	37
Days of the Week	40
Omitted Expressions	41
Evaluation	42
Conclusion	46
Appendix A: Examples of Rules	48
Appendix B: Examples of Documents	50
Bibliography	68

LIST OF TABLES

Table 1: Sources of French-language news	5
Table 2: Consistency of French and English corpora	5
Table 3: TIMEX2 tag attributes	8
Table 4: Part-of-day tokens	11
Table 5: Modifier tokens	12
Table 6: Summary of rules	29
Table 7: Results on French test set	43
Table 8: Results on French training set	43
Table 9: Results on English test set	44
Table 10: Results on English training set	44

LIST OF FIGURES

Figure 1: Learning component	16
Figure 2: Basic format of rules	17
Figure 3: System architecture	23
Figure 4: Temporal information in the French training set	25
Figure 5: Temporal information in the French test set	26
Figure 6: Temporal information in the English training set	27
Figure 7: Temporal information in the English test set	28

INTRODUCTION

In order for computers to process human language, they must have access to some sort of instructions to guide them through the language. While they obviously receive instructions in the form of software, a quite valuable source of information is annotated corpora. A corpus is defined as an “on-line collection of text and speech” (Jurafsky and Martin 194), and an annotated corpus is simply a corpus that has been marked up, or annotated, to provide additional information in a normalized fashion. Annotation is commonly used to mark parts of speech or grammatical structure, but it may also tag more complex data, such as name entities or temporal information. A corpus can be annotated either manually by hand or automatically by computer.

Automatic annotation systems have traditionally been built using human knowledge about language, but human knowledge is expensive and requires a considerable investment of time to transfer to a computer-readable format. For this reason, computational linguists have turned to corpus-based methods, which utilize a corpus to gather information about a language. Annotation rules can be learned from an annotated corpus such that they may be applied to unseen text as part of an automatic tagger. An example of a tagger that induces rules from an annotated corpus is the Brill tagger (Jurafsky and Martin 307-312) for parts of speech.

While there has been lots of research on corpus-based methods for learning annotation, little attention has been paid specifically to the learning of temporal annotation from corpora. Temporal information becomes valuable when it is correctly

annotated in corpora, thereby expanding the performance of those practical applications that make use of the corpora. In their paper “Guidelines for Annotating Temporal Information,” Mani et al. suggest that “applications that can benefit from such an annotated corpus include information extraction (e.g., normalizing temporal references for database entry), question-answering (answering ‘when’ questions), summarization (temporally ordering information), machine translation (translating and normalizing temporal references), and information visualization (viewing event chronologies)” (1). When properly annotated, temporal information can eventually “position events in time, either relatively with respect to other events or absolutely with respect to calendrical time” (Setzer and Gaizauskas 1).

Prior research in temporal annotation includes Mani and Wilson’s “Robust Temporal Processing of News” (2000) and Setzer and Gaizauskas’ “A Pilot Study on Annotating Temporal Relations in Text” (2001). Mani and Wilson’s paper describes an annotation scheme that is simple enough to promote consistency among annotators yet precise enough to capture a wide array of temporal information. It adapts ISO standards for representing points in time in calendric terms while introducing tokens to represent commonly occurring temporal units, such as *summer* and other seasons, as well as duration and time intervals. There is a compositional approach, reflecting that temporal information comes not only from time expressions alone but also from other places in text. For example, such information may come as offsets from the reference time, as in indexicals (e.g., *yesterday*, *today* and *tomorrow*), certain lexical markers

(e.g., the word *next* in *next year*) and verb tense (e.g., a future reference with future tense or a past reference with past tense).

Setzer and Gaizauskas studied the effectiveness of guidelines and technique for annotating events and the temporal relations between events in newswire texts. Their annotation method handles four types of temporal entities: events (“something that happens, something that one can imagine putting on a time map”), times (“viewed as having extent (intervals) or as being punctual (points)” and “capable of being put on a timeline”), temporal relations (“temporal relations to other events and to times”, “*before, after, includes, included* and *simultaneous*”) and event identity (to identify events when “referred to more than once, with subsequent references giving more detailed information”) (2).

In this paper, I will describe a language-independent automatic annotation system I developed that learns rules from a temporally annotated corpus in a target language. The learning process takes mere seconds for the program to complete, leaving the sole cost of the system as the cost of annotating a corpus using the TIMEX2 annotation scheme. The system relies on example-based learning whereby the annotated corpus it receives provides examples of input (i.e., tagged natural language text) and output (i.e., the tag and its attributes). The learner’s job, then, is to induce the functions, or rules, that are necessary to derive the output from the input in each example, ultimately enabling itself to derive complete TIMEX2 tags from unseen natural language text in the language of the corpus from which it learned the rules.

The following section describes in more detail this learning process with examples of real data and the nature of the rules the data provides. In that section, I will describe my approach, including corpora, methodology, and system architecture. I will then explain the results for each type of rule that the system is currently capable of generating, and finally, I will discuss evaluation results obtained by scoring the machine-annotated text against human-annotated text, using both French and English corpora.

APPROACH

Corpora

Machine learning techniques require data to be divided between a training set and a test set. The training set, which is larger than the test set, is the portion of data from which the learning component of the system induces rules. The learner never sees the test set, which is held back until the time of evaluation.

In developing my system, I focused on a corpus of French news data borrowed from the *Yahoo! France* news site on the World Wide Web in January and March of 2002. Articles on the site come from a variety of sources, including the Agence France-Presse, the Associated Press, and Reuters. A detailed breakdown of sources for the 30 documents in the French training set and the 15 documents in the French test set appears below.

<i>Source</i>	<i>Training Set</i>	<i>Test Set</i>
Agence France-Presse	13	8
Associated Press	7	4
Reuters	6	0
Other	4	3
Total	30	15

Table 1: Sources of French-language news

Both the training set and the test set were hand-annotated using the TIMEX2 scheme, which is summarized in the next subsection. Tagged documents were saved as new files, allowing the raw, untagged documents to be used later for evaluation purposes.

In total, the French training set contained 382 TIMEX2 tags, and the test set contained 204 tags, averaging 12.73 and 13.60 tags per document, respectively.

To test the portability of the system across languages, I also used a collection of temporally annotated English language news articles from January 1998’s *The New York Times*, as assembled in the TIDES Topic Detection and Tracking program (TDT-2) corpus. This corpus consists of 193 documents in total, which I split into a training set of 122 documents and a test set of 71 documents. The corpus was already annotated according to the TIMEX2 guidelines, so I simply removed the TIMEX2 tags to derive raw, unannotated data. The English training set contained a total of 2395

<i>Description</i>	<i>Tags</i>	<i>Documents</i>	<i>Tags/Document</i>
French Training Set	382	30	12.73
French Test Set	204	15	13.60
English Training Set	2395	122	19.63
English Test Set	1466	71	20.65

Table 2: Consistency of French and English corpora

TIMEX2 tags, and the test set contained 1466 tags, with an average of 19.63 and 20.65 tags per document, respectively. Clearly, the English corpus was slightly richer in

temporal information as was the French corpus. Table 2 shows a breakdown of the French and English corpora.

Methodology

By design, the system I developed does not presume to know anything about the natural language of the data on which it is trained and tested. It must therefore rely solely upon the information provided through the TIMEX2 tags themselves. Given this information, the system's learning component analyzes the portion of natural language text around which TIMEX2 tags sit, attempting to determine which elements of the natural language directly map to values within the tag, what order those elements appear in the natural language and what additional information must be learned from similar examples in the corpus. In this section, I will briefly describe the TIMEX2 annotation scheme that is such an integral part of how my system works and discuss any necessary changes to the scheme for adapting it to French. Then, I will show how the learning process works, and I will outline some limitations that exist for this approach.

Background

As Mani et al. recommend in their 2001 paper "Guidelines for Annotating Temporal Information", a temporal annotation scheme should be simple yet precise so as to be simple enough for humans to understand and execute while precise enough to capture the maximum meaning of temporal information. It should also be natural in

the distinctions it highlights in time expressions and meaning, expressive in the amount of information it can represent, and reproducible to the extent that different annotators (human or machine) would evaluate temporal information consistently (Mani et al. 2). The TIMEX2 annotation scheme they designed meets these goals. It surpasses the scheme used in the Message Understanding Conference in 1998 “not only in terms of the range of expressions that are flagged, but, also, more importantly, in terms of representing and normalizing the time *values* that are communicated by the expressions” (Mani et al. 1). Also, it can handle both fully specified and context-dependent time expressions. They point out that “more than two-thirds of time expressions in print and broadcast news were context-dependent ones” (Mani et al. 1), clearly indicating the importance of addressing context-dependent temporal information in any annotation system. Finally, they mention that their annotation scheme is flexible enough to apply to multiple genres, from news to scheduling dialogs (Mani et al. 1). For an in-depth look at the TIMEX2 scheme, please refer to the guidelines (Ferro et al. 2001) and instruction manual (Ferro 2001), summarized throughout this section.

The TIMEX2 scheme calls for the annotation of markable temporal expressions, defined as those expressions whose syntactic head contain a lexical trigger. A lexical trigger is “a word or numeric expression whose meaning conveys a temporal unit or concept” (Ferro 3) and may surface as a noun, proper noun, numeric time pattern, adjective, or adverb when used in a temporal sense. Once a temporal

expression is tagged with a <TIMEX2> tag initially and a </TIMEX2> tag finally, its meaning is captured and recorded as an attribute in the opening <TIMEX2> tag. There are seven types of attributes that may appear in a TIMEX2 tag: VAL, MOD, SET, PERIODICITY, GRANULARITY, NON_SPECIFIC and COMMENT. Table 3, reproduced from the instruction manual (Ferro 6), defines these attributes and provides an example of each.

<i>Attribute</i>	<i>Function</i>	<i>Example</i>
VAL	Contains a normalized form of the date/time	VAL="1964-10-16"
MOD	Captures temporal modifiers	MOD="APPROX"
SET	Identifies expressions denoting sets of times	SET="YES"
PERIODICITY	Captures the period between regularly recurring times	PERIODICITY="F1M"
GRANULARITY	Captures the unit of time denoted by each set member in a set of times	GRANULARITY="G3D"
NON_SPECIFIC	Identifies non-specific expressions	NON_SPECIFIC="YES"
COMMENT	Contains any comments the annotator wants to add	COMMENT="context garbled"

Table 3: TIMEX2 tag attributes

Note that, despite their use of English, the attributes of TIMEX2 tags are language-independent. That is, semantically equivalent expressions can be tagged with identical TIMEX2 tags in both English and French, as shown in examples (1) and (2) and throughout this section.

- (1) <TIMEX2 VAL="P3D">three days</TIMEX2>
<TIMEX2 VAL="P3D">trois jours</TIMEX2>
- (2) <TIMEX2 VAL="1998-08-08">August 8, 1998</TIMEX2>
<TIMEX2 VAL="1998-08-08">le 8 août 1998</TIMEX2>

Calendar dates, whether they are standalone or part of a range of dates, are annotated with an ISO-compliant value inside a VAL attribute. This is the case for absolute and fully specified dates, such as those in example (2) above as well as for relative or under-specified dates, such as those in (3) below.

(3) <TIMEX2 VAL="2002-03-29">yesterday</TIMEX2>
<TIMEX2 VAL="2002-03-29">hier</TIMEX2>

When the expression contains a time of day, that information is also reflected in the VAL attribute, as shown in (4).

(4) <TIMEX2 VAL="2002-03-30-T18:51">6:51 pm</TIMEX2>
<TIMEX2 VAL="2002-03-30-T18:51">18h51</TIMEX2>

The VAL attribute is also used for durations of time, such as *five weeks* (*cinq semaines*) or *an hour and a half* (*une heure et demie*). To represent the unit of time contained in the duration, the TIMEX2 scheme uses the ISO format shown in (5) where P is constant to each duration representation and T distinguishes times of day from larger units of time.

(5) PnYnMnDTnHnMnS

Here, Y denotes years, the first M (before the T) denotes months, D denotes days, H denotes hours, the second M (after the T) denotes minutes, and the S denotes seconds. Additional units include W (for weeks), E (for decades), C (for centuries), and L (for millennia). Examples of durations in English and French are shown in (1) above.

General references to the past, the present, or the future are assigned a VAL of PAST_REF, PRESENT_REF, or FUTURE_REF, respectively. The set of expressions

for each type of reference is more or less static, so it is relatively simple to determine the value of the VAL attribute given a text string that matches of a past, present, or future expression. Examples of PAST_REF expressions include *past*, *yesterday*, and *a long time ago* in English and *le passé (the past)* and *hier (yesterday, generally)* in French. PRESENT_REF expressions may include *now*, *current* and *these days* in English and *l'heure actuelle (the current time)*, *aujourd'hui (today, generally)*, and *actuellement (currently)* in French, and FUTURE_REF expressions may include *tomorrow* and *the future* in English and *l'avenir (the future)* and *désormais (from now on)* in French.

Subunits of years are also representable in the TIMEX2 annotation scheme. The four seasons are indicated by the tokens SP (spring), SU (summer), FA (fall), and WI (winter) in the month position of the VAL attribute, as illustrated in (6), and yearly quarters and halves are similarly annotated with the tokens Q1, Q2, Q3, Q4, H1 and H2, as shown in (7).

(6) <TIMEX2 VAL="1825-SU">the summer of 1825</TIMEX2>
<TIMEX2 VAL="1825-SU">l'été 1825</TIMEX2>

(7) <TIMEX2 VAL="2001-H2">the second half of 2001</TIMEX2>
<TIMEX2 VAL="2001-H2">deuxième semestre 2001</TIMEX2>

Likewise, subunits of days can be captured in the TIMEX2 tag with a token that represents a general time of day. These tokens, whose placement is at the hour position of the VAL attribute, are shown below in Table 4, which has been reproduced from page 15 of the instruction manual (Ferro 6).

<i>Token</i>	<i>Position</i>	<i>Expressions</i>
MO	hour	morning
MI	hour	mid-day
AF	hour	afternoon
EV	hour	evening
NI	hour	night
DT	hour	morning + afternoon (basically daytime or working hours)

Table 4: Part-of-day tokens

An example of a part-of-day token using both English and French appears in (8) below.

- (8) <TIMEX2 VAL="2002-03-25-TMO">Monday morning</TIMEX2>
<TIMEX2 VAL="2002-03-25-TMO">lundi matin</TIMEX2>

An X placeholder may be used where there is non-specificity in a temporal expression that cannot be resolved through context. In the VAL attribute, this may occur as part of a date, as in example (9), but it appears more commonly in durations, as shown in example (10).

- (9) <TIMEX2 VAL="XXXX-03-19">March 19</TIMEX2>
<TIMEX2 VAL="XXXX-03-19">19 mars</TIMEX2>
- (10) <TIMEX2 VAL="PXD">several days</TIMEX2>
<TIMEX2 VAL="PXD">plusieurs jours</TIMEX2>

Working in conjunction with the VAL attribute, the MOD attribute of the TIMEX2 tag represents modified temporal information. As the instruction manual explains, the MOD attribute is designed “to capture the basic semantics of quantifier modifiers (e.g., *approximately*, *no more than*) and lexicalized aspect markers (e.g., *early*, *start [of]*)” (Ferro 17). Table 5, which has been slightly modified from the

instruction manual (Ferro 18), outlines the available tokens for MOD, and examples (11) and (12) illustrate how they can be used.

	<i>Token</i>	<i>Sample Expressions (English)</i>
Points	BEFORE	more than (e.g., “more than a decade ago”)
	AFTER	less than (e.g., “less than a year ago”)
	ON_OR_BEFORE	no less than (e.g., “no less than a year ago”)
	ON_OR_AFTER	no more than (e.g., “no more than a year ago”)
Durations	LESS_THAN	less than (e.g., “less than 2 hours long”) nearly (e.g., “nearly four decades”)
	MORE_THAN	more than (e.g., “more than 5 minutes”)
	EQUAL_OR_LESS	no more than (e.g., “no more than 10 days”)
	EQUAL_OR_MORE	at least (e.g., “at least 10 days”)
Points and Durations	START	early (e.g., “the early 1960s”) dawn (e.g., “the dawn of 2000”) start (e.g., “the start of the quarter”) beginning (e.g., “the beginning of the year”)
	MID	middle (e.g., “the middle of the month”) mid- (e.g., “mid-February”)
	END	end (e.g., “the end of the season”) late (e.g., “late June”)
	APPROX	about (e.g., “about three years ago”) around (e.g., “around Christmas”)

Table 5: Modifier tokens

(11) <TIMEX2 VAL=“2002” MOD=“END”>late 2002</TIMEX2>
<TIMEX2 VAL=“2002” MOD=“END”>fin 2002</TIMEX2>

(12) <TIMEX2 VAL=“P65Y” MOD=“MORE_THAN”>more than 65
years</TIMEX2>
<TIMEX2 VAL=“P65Y” MOD=“MORE_THAN”>plus de 65
ans</TIMEX2>

In addition to points of time and durations of time, TIMEX2 handles set-denoting temporal expressions, which are expressions that tell how often something happens. To accomplish this, one makes use of the SET, GRANULARITY, and PERIODICITY tag attributes. The SET attribute simply states that the tagged

expression denotes a set. The GRANULARITY attribute encapsulates the unit of time denoted by each set member, and the PERIODICITY attribute expresses the length of time between each set member. Both the GRANULARITY and PERIODICITY tag attributes take the format of a VAL duration, replacing the initial “P” with a “G” or “F”, respectively. Examples (13) and (14) demonstrate how set-denoting temporal expressions are tagged with the TIMEX2 annotation scheme.

- (13) <TIMEX2 SET=“YES” GRANULARITY=“G1Y”
PERIODICITY=“F1Y”>each year</TIMEX2>
<TIMEX2 SET=“YES” GRANULARITY=“G1Y”
PERIODICITY=“F1Y”>chaque année</TIMEX2>
- (14) <TIMEX2 SET=“YES” GRANULARITY=“G1D”
PERIODICITY=“F1D”>daily</TIMEX2>
<TIMEX2 SET=“YES” GRANULARITY=“G1D”
PERIODICITY=“F1D”>quotidienne</TIMEX2>

Like the past, present, and future references discussed earlier, set-denoting expressions tend to be fairly consistent in a language. In other words, the English expression *daily*, used above in example (14), will always be tagged as a set with both a granularity and a periodicity of one day.

The remaining TIMEX2 tag attributes – NON_SPECIFIC and COMMENT – occur less frequently than the other attributes. While the COMMENT attribute is applicable whenever the annotator wishes to comment on the tagged expression, the NON_SPECIFIC attribute is used for lexical triggers used generically, singular indefinite lexical triggers, and various other non-specific temporal expressions. The attribute takes only “YES” as its value, as shown in (15).

(15) <TIMEX2 NON_SPECIFIC="YES">a decisive year</TIMEX2>
<TIMEX2 NON_SPECIFIC="YES">une année décisive</TIMEX2>

Note that the NON_SPECIFIC attribute may appear alongside a VAL attribute with the X placeholder. This is useful in cases when there is some temporal information to be accounted for, even if the expression as a whole is non-specific.

Some temporal expressions require only the TIMEX2 tags themselves without attributes. This happens with event-anchored temporal expressions such as *five minutes after the announcement*. Finally, coreferring temporal expressions should be annotated exactly as the expressions with which they corefer. That is, if two expressions corefer, the greater specificity of one carries over to the other.

The meaning of temporal expressions captured within the TIMEX2 tag is not the only factor an annotator must consider. The beginning and end tags must surround the correct extent of the expression, as explained in the instruction manual (Ferro 29-34). A properly tagged expression belongs to one of the following grammatical categories: noun, noun phrase, adjective, adverb, or adjective/adverb phrase. Syntactically, the expression may include premodifiers, such as determiners, possessive pronouns, and modifiers whose meaning could be captured in a MOD attribute (see Table 5), and postmodifiers, like prepositional phrases and dependent clauses. In the case of appositives, range expressions, conjoined expressions and embedded expressions, individual expressions should be tagged separately. Lastly, multiple expressions can be tagged singly if they appear together and if they express

the same value (e.g., *twelve noon*) or express values that are hierarchically related (e.g., *Friday evening*).

The TIMEX2 scheme, which was designed for English, carries over to French quite well, which is not too surprising considering that English and French are relatively similar languages. While no changes were necessary concerning the values written to TIMEX2 attributes, the extent of tags differs between English and French. One example is the TIMEX2 scheme's requirement that an entire noun phrase should be included within the tags when the noun head of the phrase is a markable lexical trigger. That is, if an English text contains the lexical trigger *past* preceded by the determiner *the*, it would be tagged as shown in (16).

(16) <TIMEX2 VAL="PAST_REF">the past</TIMEX2>

The same is true of the equivalent French expression *le passé*, shown in (17).

(17) <TIMEX2 VAL="PAST_REF">le passé</TIMEX2>

In French, however, when certain prepositions (e.g., *de* (*of, from*), *à* (*to*)) precede the determiner *le*, the preposition and determiner merge to form a partitive article (e.g., *du, au*). I opted to not include this article within TIMEX2 tags, so the expression *du passé* would be tagged as shown in (18).

(18) du <TIMEX2 VAL="PAST_REF">passé</TIMEX2>

Learning Component

In the previous subsection, I described the format of the TIMEX2 tags and the type of information they may contain. It is with this knowledge that the learning

component of my system teaches itself about the natural language of the corpus it reads as input. The first major step of the learner is to extract TIMEX2 tags and the natural language expressions they annotate from a training set. Next, it groups the tagged expressions according to their type (i.e., their TIMEX2 attributes) and specificity (i.e., the degree of information captured in the TIMEX2 tags). Within each group, the learner looks to the natural language expressions to find information that may relate to information in the tags, ultimately inventing rules that translate natural language text of a determined pattern to attributes in a TIMEX2 tag. Figure 1 illustrates the basic flow of this learning method. The four steps at the bottom of the illustration, ending in writing the rule, repeat for as many rules the system is designed to support.

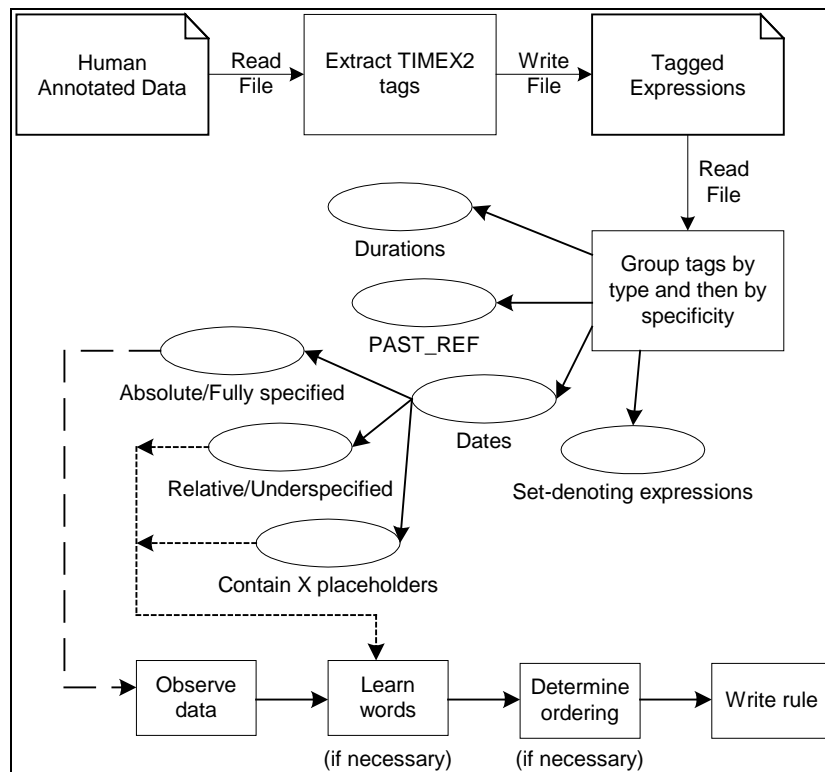


Figure 1: Learning component

Note that the focus of the “Observe data” step is fully specified data. The reason for this is that the learner must use the standard form of the natural language that contains enough explicit information to populate the TIMEX2 tag. Only after it has gained this information as its foundation, it can attempt to make use of less specified examples.

Each rule shares a common structure, and the learner fills in the gaps in the structure to construct the final working rule. Figure 2 shows the skeleton form of a rule. Items enclosed in angle brackets vary rule to rule.

<p><i>Rule Name:</i> <Name></p> <p><u>Setup</u></p> <p><i>Input String:</i> Untagged document (untagged_string)</p> <p><i>Variables:</i> altered, more_to_read, tagged_string <Additional Variables></p> <p><u>Text Processing</u></p> <p><i>Pattern:</i> <NL Pattern to Match (regular expression)></p> <p><i>Process:</i> <Instructions (map value or evaluate expression)></p> <p>If <i>Pattern</i> found in <i>Input String</i></p> <ul style="list-style-type: none">Perform <i>Process</i> on matching substringWrap matching substring with TIMEX2 tagsAssign result of <i>Process</i> to TIMEX2 attributesWrite altered string to tagged_string <p><u>Closing</u></p> <p>If <i>Input String</i> has been altered, output tagged_string. Otherwise, output the <i>Input String</i>.</p>

Figure 2: Basic format of rules

See Appendix A for more detailed pseudocode examples of rules.

Let’s consider an example. Example (19) contains a list of tagged expressions from the French training set. These expressions have been isolated because of their similarity of type and specificity, all annotated with a digital year, month, and date in the VAL attribute and including at least three words of natural language, one of which

is a direct match with the year element of the VAL attribute and another which contains a match with the date element of the VAL attribute.

- (19) <TIMEX2 VAL="2001-11-20">20 novembre 2001</TIMEX2>
<TIMEX2 VAL="2005-01-01">1er janvier 2005</TIMEX2>
<TIMEX2 VAL="2002-01-01">1er janvier 2002</TIMEX2>
<TIMEX2 VAL="2001-12-31">le 31 décembre 2001</TIMEX2>
<TIMEX2 VAL="2002-01-01">le 1er janvier 2002</TIMEX2>
<TIMEX2 VAL="2001-12-31">31 décembre 2001</TIMEX2>
<TIMEX2 VAL="2002-01-06">6 janvier 2002</TIMEX2>
<TIMEX2 VAL="2000-05-18">le 18 mai 2000</TIMEX2>
<TIMEX2 VAL="1962-03-19">le 19 mars 1962</TIMEX2>
<TIMEX2 VAL="1954-11-01">1er novembre 1954</TIMEX2>
<TIMEX2 VAL="1962-07-02">le 2 juillet 1962</TIMEX2>
<TIMEX2 VAL="1999-06-10">le 10 juin 1999</TIMEX2>
<TIMEX2 VAL="2002-06-27">27 juin 2002</TIMEX2>
<TIMEX2 VAL="2002-01-15">15 janvier 2002</TIMEX2>
<TIMEX2 VAL="2001-02-17">le 17 février 2001</TIMEX2>
<TIMEX2 VAL="1997-08-03">le 3 août 1997</TIMEX2>
<TIMEX2 VAL="1999-04-30">le 30 avril 1999</TIMEX2>
<TIMEX2 VAL="1996-06-27">le 27 juin 1996</TIMEX2>
<TIMEX2 VAL="1997-07-04">le 4 juillet 1997</TIMEX2>
<TIMEX2 VAL="2001-01-15">le 15 janvier 2001</TIMEX2>

To simplify, the learner must focus on those tags with the minimal number of words that still contain all the information represented in the TIMEX2 tag. Because the VAL attribute holds three pieces of information (year, month, and date), it reduces the list to just three words in the natural language text.¹ The resulting list, which filters out all the occurrences with more than three words, is shown in (20).

- (20) <TIMEX2 VAL="2001-11-20">20 novembre 2001</TIMEX2>
<TIMEX2 VAL="2005-01-01">1er janvier 2005</TIMEX2>
<TIMEX2 VAL="2002-01-01">1er janvier 2002</TIMEX2>
<TIMEX2 VAL="2001-12-31">31 décembre 2001</TIMEX2>

¹ This is a limitation to my approach, which assumes that each month name in the natural language consists of a single word. To accommodate languages with multi-word month names, the learner would need to lift this assumption and look for potential month names of n number of words.

<TIMEX2 VAL="2002-01-06">6 janvier 2002</TIMEX2>
<TIMEX2 VAL="1954-11-01">1er novembre 1954</TIMEX2>
<TIMEX2 VAL="2002-06-27">27 juin 2002</TIMEX2>
<TIMEX2 VAL="2002-01-15">15 janvier 2002</TIMEX2>

Using this simplified list, the learner can determine the position of each element of the date in the tagged natural language text by comparing it to the year, month, and date values provided in the VAL attribute. In each of the tags in (20), the learner observes that the element matching VAL's date appears first, and the element directly matching VAL's year appears third. By process of elimination, the month element must be in the second position. This gives us the rule that, in the language of the input corpus, DD Month YYYY corresponds to YYYY-MM-DD in the TIMEX2 tag's VAL attribute. DD Month YYYY roughly corresponds to the *Pattern* piece of the rule, while the mapping corresponds to the *Process* piece of the rule, as outlined above in Figure 2. Now, the system can return to the four-word date examples that were filtered out between example (19) and example (20). Because the final three words of the text contain all the temporal information needed to write the TIMEX2 tag correctly, the initial word – *le* in this case – does not add meaning to the expression. The learner concludes that this extra word must be a determiner and writes into the *Pattern* piece of the rule to include that word within the tag if it appears before the date pattern.

Pseudocode for the full date rule appears in Appendix A at the end of this paper.

The system can learn other patterns through this same general methodology. The learner isolates similar TIMEX2 tags, identifies commonalities among the text they surround, and writes rules to capture the link between natural language text and

the TIMEX2 tag and its attributes. However, because the system knows nothing about the language of the corpus, an important task of the learner throughout this process is to learn the meaning of certain words. For example, the expressions in (19) and (20) clearly show that the names of the months are important to writing the correct two-digit month to the VAL attribute. This task can only be accomplished if the tagger can recognize the names of the months, which must be listed in the rule's *Pattern* section. The rule generator provides this information to the rule after it has learned the names of the months from the training data. There are several types of tags the learner can use to do this. The most effective are those whose VAL attribute contains a month element and whose tagged text includes a single word. These examples are the best to begin with because they contain the minimal amount of natural language necessary to identify the month reflected in VAL's month element. A list of these expressions from the French training set appears in (21).

- (21) <TIMEX2 VAL="2001-09">septembre</TIMEX2>
<TIMEX2 VAL="2003-03">mars</TIMEX2>
<TIMEX2 VAL="2001-12">décembre</TIMEX2>
<TIMEX2 VAL="2001-12">décembre</TIMEX2>
<TIMEX2 VAL="2002-01">janvier</TIMEX2>
<TIMEX2 VAL="2002-01">janvier</TIMEX2>
<TIMEX2 VAL="2002-01">janvier</TIMEX2>
<TIMEX2 VAL="2002-02">février</TIMEX2>
<TIMEX2 VAL="2001-09">septembre</TIMEX2>
<TIMEX2 VAL="2002-03">mars</TIMEX2>
<TIMEX2 VAL="2002-06">juin</TIMEX2>
<TIMEX2 VAL="2002-02">février</TIMEX2>
<TIMEX2 VAL="2002-04">avril</TIMEX2>
<TIMEX2 VAL="2002-05">mai</TIMEX2>

With this information, the learner can accurately deduce the words for eight of the twelve months of the year. It writes what it has learned to a table with one column for the names of the months and another column for the corresponding two-digit month representations used in the TIMEX2 tag. To fill the remaining four rows of the table, it can now look to slightly more complicated examples, including those with VAL to the month position and a word and year in the tagged text and full dates like those in (19) and (20). It is, of course, possible that the learner will incorrectly learn the name of a month. For example, using the French training set, the learner incorrectly learned the tenth month of the year as *dernier* (*last*) rather than *octobre*. To correct this, it must validate what it believes to be the names of months once all twelve rows of the table have been filled. For each of its month names, it cycles through every tagged expression in the training set. If each example in which the candidate month name appears also contains what the system believes the associated two-digit TIMEX2 month to be, the month is good. Otherwise, if the candidate month name appears in expressions that are not tagged with the corresponding month in the VAL tag, the month name must be incorrect, and the learner tries another word from the examples for that month. In the case of the learner incorrectly learning the tenth month in French, the learner's validation step isolated all of the examples from the training set in which the tenth month was represented, as shown in (22).

(22) 16 octobre
le 7 octobre dernier
octobre dernier
le 17 octobre

As we can see from these examples, the first pick word *dernier* only appears in two out of the four examples, plus it appears in examples that are not annotated as being in month 10. Excluding numeric words, the learner must choose between the candidate words *le (the)* and *octobre (October)*. Like *dernier*, *le* appears in just two of four examples plus other examples not tagged as month 10, but *octobre* appears in all four and only in those four, flagging it as the best choice word for the tenth month of the year. In tests with both French and English, the learner correctly validated incorrect months, ultimately learning the correct names of all twelve months of the year.

Limitations

The learning methodology of this system exhibits several limitations that may hinder its performance on certain other languages. First, it takes for granted that the corpora are encoded in the ASCII character set. Second, it assumes that the words of the natural language text are segmented by spaces. This would require pre-processing for texts in languages whose words are not yet segmented. Third, languages with word ordering that drastically differs from French or English may experience problems. The system is designed to be flexible with word ordering, but it cannot anticipate every possible ordering in a language. An example of this problem occurring even within English has to do with full dates. Using the English test set, the learner determines that the name of a month begins a fully specified date, followed by the day and year. Occasionally, however, the month and day are reversed, as in a British English news article. Trained on the English corpora used in this experiment, such dates would not

be recognized as full dates. Fourth, extent of time expressions varies across languages, as we witnessed above while adapting the TIMEX2 scheme to French. While the system may pick up on some variations, others may be missed. The issue of extent also relates to the potential problem noted above wherein a language uses multiple words to express what is represented by a single word in English or French, such as the names of months or days of the week.

System Architecture

Figure 3 illustrates the architecture and general flow of the system.

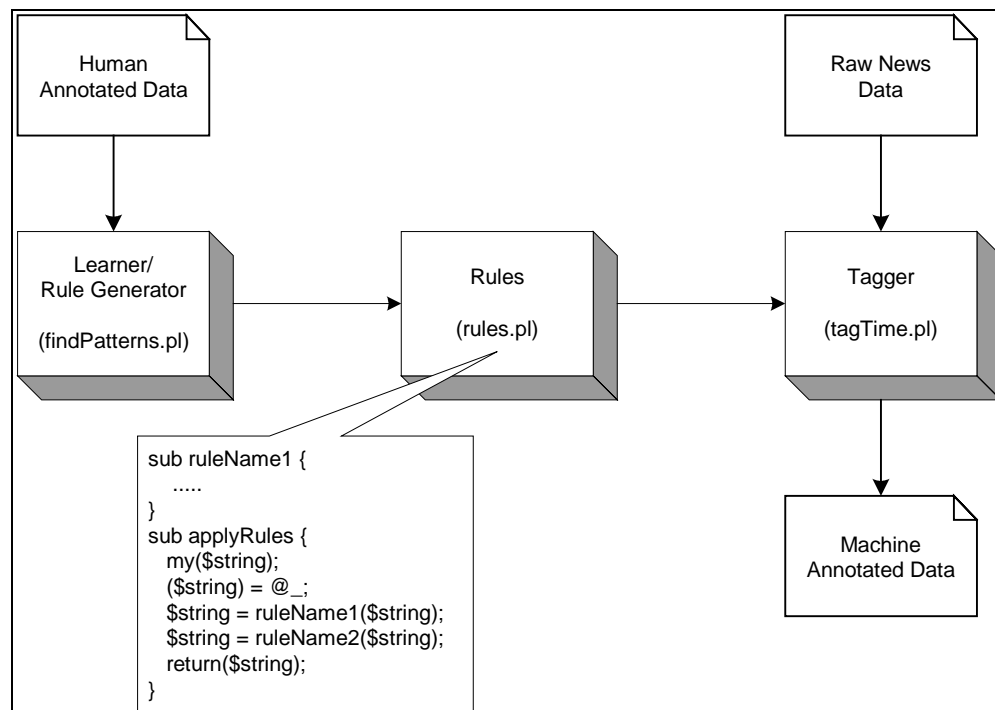


Figure 3: System architecture

The input for the Learner/Rule Generator is human-annotated data in any language, following the guidelines of the TIMEX2 scheme, explained above. The

learner reads the data and generates rules, which it writes to a separate file, represented in the diagram as Rules. The rules file consists only of Perl subroutines for each rule learned from the training data. Because the list of rules varies according to the training data, the rules file contains a master control subroutine that runs each of the rule subroutines within the rules file. The Tagger, which takes raw data in the language of the training set as input, includes the rules file so that it can reference the subroutines within it through the master control subroutine, called `applyRules`. The resulting text string from the tagger's call to the rules file is temporally annotated data, which it outputs as a new file.

Naturally, the bulk of the system lies in the Learner/Rule Generator, which must assemble all the necessary information for the rules file to be written and used correctly. The tagger itself is extremely simple, relying upon text processing in the rules file. Please refer to Appendix B at the end of this paper for sample human- and machine-annotated documents from French and English data sets.

RESULTS

Rules

While my system does not yet learn rules for all possible types of information that may appear in a TIMEX2 tag, it does handle a sampling of rules that is substantial enough to illustrate the potential of the approach. Figures 4, 5, 6, and 7 itemize the

types of temporal information contained in the French and English training and test sets.




Figure 4: Temporal information in the French training set

While it is interesting to note the variance of each type of expression across all the data sets, the most frequently occurring types of temporal information in each set remain

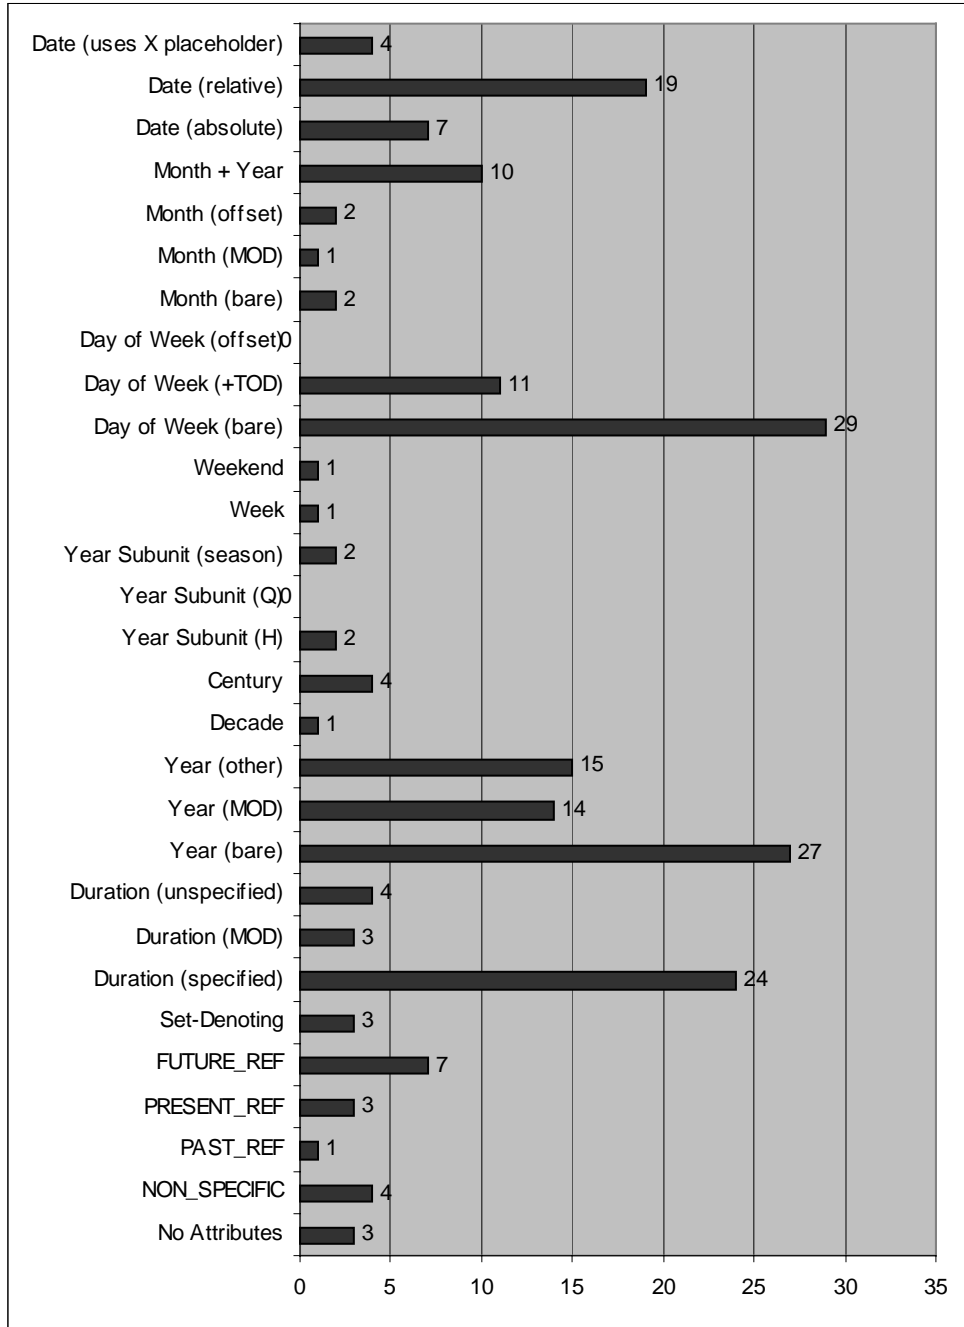


Figure 5: Temporal information in the French test set

fairly consistent: relative dates, bare days of the week, absolute years, relative years, and specified durations.




Figure 6: Temporal information in the English training set

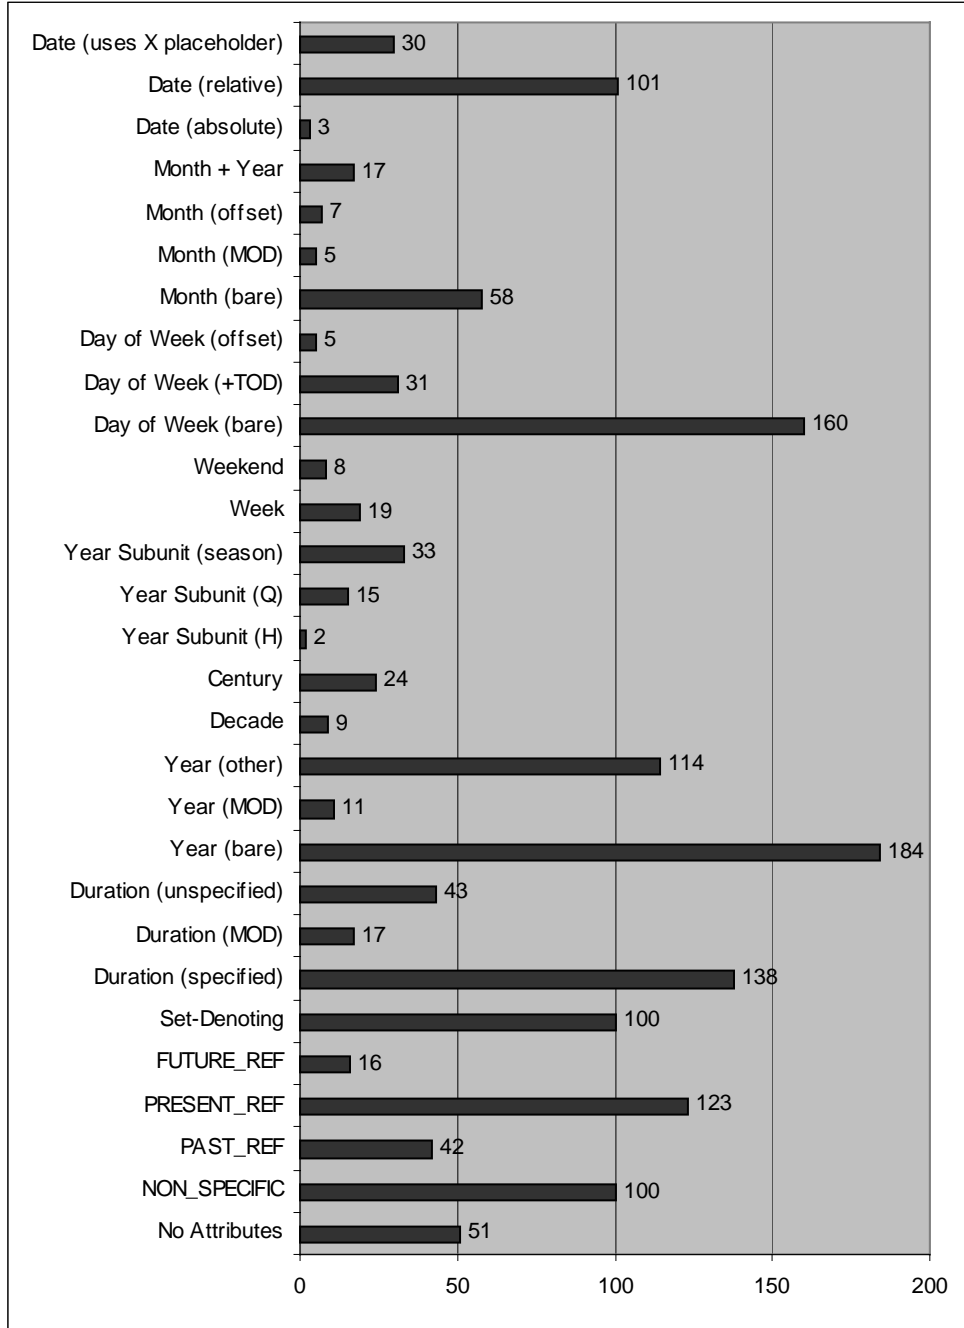


Figure 7: Temporal information in the English test set

A couple of peculiar dissimilarities between the French and English corpora are the greater relative frequency of PRESENT_REF expressions in English (8.42%) versus

French (1.71%) and the scarcity of absolute, fully specified dates in the English data (0.21%) compared to the French data (4.78%).

As one may expect, some rules are more straightforward than others, but ironically, some of those rules that present more challenges for the learner are ultimately more accurate and robust than the simple rules. The overall set of rules fall into two distinct categories: those learned and used via memoization and those learned via instance-based methods. Memoization, a timesaving technique by which input/output pairs are cached for quick retrieval, is the basis for the PAST_REF, PRESENT_REF, FUTURE_REF, and frequency rules. The remaining rules rely on instance-based learning, whereby the learner developed generalizations based upon examples in the training data. Table 6 provides a high-level summary of rules covered in this section.

<i>Method</i>	<i>Type</i>	<i>Rule</i>
Memoization	PAST_REF	$[W_1 \dots W_n] \rightarrow \text{PAST_REF}$
	PRESENT_REF	$[W_1 \dots W_n] \rightarrow \text{PRESENT_REF}$
	FUTURE_REF	$[W_1 \dots W_n] \rightarrow \text{FUTURE_REF}$
	Frequencies (set-denoting expressions)	$[W_1 \dots W_n] \rightarrow \text{SET} = X(\text{cached}/\text{YES}) +$ $\text{GRANULARITY} = Y(\text{cached}) +$ $\text{PERIODICITY} = Z(\text{cached})$
Instance- based	Durations	$[W_1 W_2] \rightarrow \text{Num}(W_1) \text{Unit}(W_2)$
	Dates	$[W_1] \rightarrow \text{YYYY}(W_1)$ $[W_1 W_2] \rightarrow \text{YYYY}(W_2) \text{MM}(W_1)$ $[W_1] \rightarrow \text{YYYY}(\text{XXXX}) \text{MM}(W_1)$ $[W_1 W_2 W_3] \rightarrow \text{YYYY}(W_3) \text{MM}(W_2) \text{DD}(W_1)$ $[W_1 W_2 W_3 W_4] \rightarrow W_1 \text{YYYY}(W_4) \text{MM}(W_3)$ $\text{DD}(W_2)$
	Days of the week	$[W_1] \rightarrow \text{YYYY-MM-DD}(W_1)$ $[W_1 W_2] \rightarrow \text{YYYY-MM-DD}(W_1) \text{Time}(W_2)$

Table 6: Summary of rules

PAST_REF, PRESENT_REF, and FUTURE_REF

Within a language, the expressions we use to refer to the past, present, or future tend to be rather static. The tagger needs only to recognize these expressions, tag them, and write the appropriate value (e.g., PAST_REF, PRESENT_REF, or FUTURE_REF) to the VAL attribute of the opening tag. To do this, the rules must provide the tagger with a list of qualifying expressions for each type. Examples (23), (24), and (25) list the French expressions that were tagged with PAST_REF, PRESENT_REF, and FUTURE_REF values, respectively.

(23)	il y a 65 millions d'années	(65 million years ago)
	il y a 170.000 ans	(170,000 years ago)
	a saison dernière	(last season)
	passé	(past)
	hier	(yesterday)
	le passé	(the past)
(24)	l'heure actuelle	(the current time)
	d'aujourd'hui	(of today)
	aujourd'hui	(today)
	actuellement	(currently)
(25)	l'avenir	(the future)
	désormais	(from now on)
	avenir	(future)

With this information, the learner writes a rule called `pastRef` that looks for strings that match the expressions in (23). When it finds a match, it simply rewrites it with “<TIMEX2 VAL=“PAST_REF”>” preceding the match and “</TIMEX2>” trailing the match. Likewise, it writes rules called `presentRef` for the expressions in (24) and a rule called `futureRef` for the expressions in (25), writing the appropriate value to the

VAL attribute for the expressions it finds. When trained on data in any other language, the rules are written to look for whatever past, present, and future expressions that were annotated in the training set for that language.

This method proved to work quite well for most past, present, and future references. However, given the relatively small sizes of the test sets, it is not surprising that the learner overfit the list of qualifying expressions to the training set, picking up rare expressions in the training data that do not appear as past, present, and future references in the test data. Examples of some of these rare expressions that were tagged in the human-annotated English training set appear in (26), (27), and (28) below for past, present, and future.

(26) his communist past
earlier, less virulent days
the early days of sound movies
that point in time

(27) the time to move
more than 17 months later

(28) the future of American politics
the financial and legal future of the tobacco industry

A commonality among these expressions is their wordiness relative to the more common past, present, and future references, like *past*, *yesterday*, *now*, *current*, and *the future*. In some of the rare examples, particularly those in (28) which are both variations on the expression *the future*, knowledge of the language via part-of-speech tags or parsed text could have been beneficial in creating rules that are more flexible to adjectives and postmodifiers within the expression. However, such knowledge is not

available to the learner or the tagger in my system, consequently allowing variations of past, present, and future expressions to go unnoticed or to be only partially tagged. Precisely this occurred both with my French and English test sets, incorrectly missing postmodifiers in the extent of the TIMEX2 tags. Examples of incorrectly tagged future expressions in my French and English test sets appears in (29) and (30).

- (29) Expression: l'avenir de la nation française toute entière
(the future of the entire French nation)
Tagged by system as: <TIMEX2 VAL="FUTURE_REF">l'avenir</TIMEX2> de la nation française toute entière
Tagged by human as: <TIMEX2 VAL="FUTURE_REF">l'avenir de la nation française toute entière</TIMEX2>
- (30) Expression: the future of managed health care
Tagged by system as: <TIMEX2 VAL="FUTURE_REF">the future</TIMEX2> of managed health care
Tagged by human as: <TIMEX2 VAL="FUTURE_REF">the future of managed health care</TIMEX2>

Somewhat surprisingly, human annotators did not always outperform the machine in tagging past, present, and future references. By nature, a computer is more consistent than a human in carrying out repetitive tasks. The scorer pinpointed numerous occurrences in the test sets of valid past, present, or future references – particularly markable adjectives like the English words *current* and *former* – that were correctly tagged by the machine yet overlooked by the human annotator.

Frequencies

The rule for annotating frequencies is very similar to the rules for annotating past, present, and future references in that it scans text for matches to examples of set-

denoting expressions found in the training data. Some of the expressions found in the French and English training sets are listed in (31) and (32) below along with their corresponding values for the SET, GRANULARITY, and PERIODICITY attributes.

(31) tous les ans (every year)	YES	G1Y	F1Y
par an (per year)	YES	G1Y	F1Y
chaque année (each year)	YES	G1Y	F1Y
tous les jours (every day)	YES	G1D	F1D
par mois (per month)	YES	G1M	F1M
quotidienne (daily)	YES	G1D	F1D
chaque décennie (each decade)	YES	G1E	F1E
(32) each year	YES	G1Y	F1Y
many hours a day	YES	GTXH	F1D
every night	YES	G1NI	F1D
weekly	YES	G1W	F1W
a year	YES	G1Y	F1Y
annually	YES	G1Y	F1Y
every five years	YES	G1Y	F5Y

The learner writes a rule called setExp that looks for strings that match the set-denoting expressions found in the training set. When it finds a match, it simply rewrites it with preceding and trailing TIMEX2 tags, filling in the SET, GRANULARITY, and PERIODICITY attributes according to their values from the training set examples. When trained on data in any other language, the rules are written to look for the set-denoting expressions that were annotated in the training set for that language.

Unfortunately, while the French test set contained three set-denoting expressions, none of the twelve set-denoting expressions in the French training set reappeared in the French test set. The English corpus, being significantly larger, did have some overlap in set-denoting expressions across both the training set and the test

set. Clearly, a large training set is an advantage for any learning task, as demonstrated by the failure of this hard-coding technique for set-denoting expressions in my French corpus. If the French training set had been larger, it might have included some of those expressions contained in the test set.

A future version of this language-independent temporal annotation system should build a lexicon of words that occur in set-denoting expressions and recognize the patterns in which they appear in text. For example, set-denoting expressions in French often begin with *tous les* (*every*), *chaque* (*each*), or *par* (*per*), followed by a word that – as learned from examples – represents some period of time. The same can be said for the English words *each*, *every*, and even *a*. Having learned this information, the system could generate a more flexible rule for sets, one that seeks a pattern, not a particular string.

Durations

In order to recognize durations of time, the tagger needs to know the words of the target language that represent temporal units. In French, such words include *ans* (*years*), *mois* (*months*), *semaines* (*weeks*), and *jours* (*days*), and in English, they include *years*, *months*, *weeks*, and *days*, to name a few. The learner, of course, must build this knowledge and write it into the duration rule for the tagger to reference. To learn these words, the learner isolates the tags in the training data whose values represent durations (see (5) above). It divides these examples into subgroups based on the temporal unit they contain. For each tag it finds for the temporal unit *Y*, it writes

each non-numeric word to a table of candidate words for *year* that stores the candidate words and the number of instances for the temporal unit *Y* in which that word was used. Example (33) lists the candidate words and their frequencies for *year* in the French training set.

(33)	quarante	1
	ans	15
	nombreuses	1
	les	1
	dernières	2
	plus	3
	dix	2
	pour	1
	vingt	1
	de	6
	parité	2
	cinq	2
	plusieurs	1
	prochaines	1
	fixe	1
	d'un	1
	nombreuses,	1
	100.000	1
	500.000	1
	peso	1
	précédentes	1
	ces	1
	des	1
	sept	1
	années	8
	dollar	1

After looping through all qualifying examples, it selects the word from the table of candidate words with the highest frequency value and deems it to be the most likely word in the target language for *year*. In the French example shown in (33), the learner

correctly selects the word *ans* to represent *years* in French. This process repeats for training examples of other temporal units.

The rule finds occurrences of these temporal unit words. When it finds a match, it looks at the word immediately preceding the temporal unit word. If the word consists of nothing but numerals, it annotates the two words with the VAL attribute, using the numeric value of the initial word and the meaning of the final word to complete the value of the VAL attribute. If, however, the word preceding the temporal word does not consist of numerals, it annotates the two words as a duration of an unspecified number of units.²

This overall method correctly annotates many durations of time, but it also introduces several problems. First, when a person's age appears in French news (e.g., *John Smith, 68 ans*), the tagger recognizes it as a duration of time and incorrectly tags it as such. Ages in English news, however, are expressed only by their number (e.g., *John Smith, 68*), and therefore are not recognized by the tagger. Second, the TIMEX2 annotation scheme calls for postmodifiers to be included within the extent of a tag. It is common for a temporal duration to have a postmodifier, as occurs quite frequently in my French test set, but my annotation system fails to recognize postmodifiers, instead closing the tag after the temporal unit word. This is comparable to the extent problem noted above for past, present, and future references.

² Currently, the learner does not learn the spelled-out versions of numbers, but this is a feature that it can incorporate into a future version of the program. For now, spelled-out numbers in duration expressions are simply given the value *X* in the VAL attribute.

Full and Partial Dates

In the Methodology subsection of the previous section, I described in detail the process the learner uses for full and partial dates. To briefly recap, it isolates fully specified dates and learns from them the order of the year, month, and date units as they appear in the natural language expressions. It also learns and validates the twelve months of the year and writes this information into any rule that would need it, and it takes any extraneous tagged words preceding fully specified dates to be determiners.

Partial dates are those that are missing some component of the fully specified date, such as a year, month, or date. Examples of partial dates in the French corpus include those listed below in (34) through (39).

- (34) <TIMEX2 VAL="1999-08">août 1999</TIMEX2>
<TIMEX2 VAL="2000-01">janvier 2000</TIMEX2>
<TIMEX2 VAL="2001-10">octobre 2001</TIMEX2>
- (35) <TIMEX2 VAL="2002-06">juin</TIMEX2>
<TIMEX2 VAL="2002-01">janvier</TIMEX2>
<TIMEX2 VAL="2001-12">décembre</TIMEX2>
- (36) <TIMEX2 VAL="2001-12-20">le 20 décembre</TIMEX2>
<TIMEX2 VAL="2002-01-31">31 janvier</TIMEX2>
<TIMEX2 VAL="2002-02-19">le 19 février</TIMEX2>
- (37) <TIMEX2 VAL="2002-03-22">22</TIMEX2>
<TIMEX2 VAL="2002-04-29">le 29</TIMEX2>
<TIMEX2 VAL="2002-05-08">le 8</TIMEX2>
- (38) <TIMEX2 VAL="1998">1998</TIMEX2>
<TIMEX2 VAL="2002">2002</TIMEX2>
<TIMEX2 VAL="2001">2001</TIMEX2>
- (39) <TIMEX2 VAL="1999">99</TIMEX2>
<TIMEX2 VAL="1996">96</TIMEX2>

These examples represent varying levels of difficulty for the learner and tagger. The examples in (34), for instance, are simple once the system has learned the months of the year, as they require no information outside of what it explicit in the tagged text. The same is true of the examples in (38), in which the value assigned to the VAL attribute is exactly equal to the numerals in the tagged expression.

The other examples are more complex and therefore more prone to error. For instance, the examples of bare months in (35) provide enough information to determine the month element of the VAL attribute, but they alone give no indication of the value of the year element. Similarly, the examples in (36) provide information on the month and date elements, but nothing on the year element. Both sets of examples are instances of context-dependent expressions. The TIMEX2 guidelines define context-dependence as occurring when “the intended value of a time expression may require an inference based on information elsewhere in the sentence (local context dependence) or information in the surrounding discourse (global context dependence)” (Ferro et al. 3). One method of dealing with context-dependence is to rely on coreference by which the context-dependent expression likely relates to the immediately preceding temporal reference. For now, my system tags these particular types of context-dependent expressions with placeholder Xs in the place of the unknown year. The examples in (37) and (39) are even more dependent on context than are those in (35) and (36). Here, only one element of the full date is provided, leaving even more to context and

widening the opportunity for mistakes. My system makes no attempt to annotate such examples.

So far, there has been no reason for rule ordering, but full and partial dates necessitate it to avoid multiple tags where only one is needed. Without rule ordering, a fully specified French date like

20 novembre 2001

could erroneously be tagged as

```
<TIMEX2 VAL="2001-11-20">20 <TIMEX2 VAL="2001-11"><TIMEX2  
  VAL="2001-11">novembre</TIMEX2> <TIMEX2  
  VAL="2001">2001</TIMEX2></TIMEX2></TIMEX2>
```

To force rule ordering, the learner writes rules pertaining to full or partial dates to the master control loop of the rules file in such an order that more specified dates are tagged before less specified dates and rules for less specified dates check for nearby TIMEX2 tags before inserting additional tags around text.

The system handles full and partial dates fairly well, but yet again, the issue of extent surfaces as a problem that needs to be addressed in a future version. Particularly with context-dependent expressions like those in (35) and (36), the words *prochain* (*next*) and *dernier* (*last*), which may appear alongside a temporal expression, should actually be included within the tag and ideally utilized to resolve the value of the year element for the VAL attribute of the TIMEX2 tag. The same is true of the words *next* and *last* in English.

Days of the Week

Bare days of the week are among the most frequently occurring time expressions in news text, accounting for 13.65% of tagged expressions in the French corpus and 7.77% of tagged expressions in the English corpus. At first glance, these context-dependent expressions seem likely to lead to ambiguity in resolving their values for the TIMEX2 tag.

To learn the days of the week, the learner loops through each example in the training set in which a single non-numeric word is given a full date value in the VAL attribute, isolating the tagged words as names of the days of the week and ordering them based on the date information given in the VAL attributes. With this information from the learner for the most frequently occurring year in the training set, the tagger can determine the name for any other day of that year and, in turn, for any day of any other year. An analysis of the corpora shows that 96.20% of day-of-week expressions in the French corpus and 96.97% of day-of-week expressions in the English corpus refer to dates within a seven-day window around the document date whereby the document date falls in the middle of the window. The rule for annotating days of the week builds this window and uses it to map days of the week in text to a calendar date for the VAL attribute. For example, if a document's date were January 5, 1998 in the English test set, the seven-day window would look like (40) below.

- (40) “Friday” → 1998-01-02
“Saturday” → 1998-01-03
“Sunday” → 1998-01-04
“Monday” → 1998-01-05 (document date)
“Tuesday” → 1998-01-06
“Wednesday” → 1998-01-07
“Thursday” → 1998-01-08

The tagger’s task would then be simple. Where it recognizes the word “Monday” it gives it a VAL of 1998-01-05, a “Tuesday” is given 1998-01-06, and so on.

Along with bare days of the week are temporal expressions that consist of a day of the week plus a time of day, as represented in the VAL attribute as described earlier in Table 4. After having established a successful method for deriving the value for a VAL attribute from names of days of the week, learning the times of the day is simple. The learner isolates the tagged expressions in the training set that include each of the tokens from Table 4. Ignoring any tagged words that were determined to be names of days of the week, the learner ascertains the word that corresponds to each time of day, writing these words into the rule to append the correct time-of-day token onto the value in the VAL attribute.

Omitted Expressions

For lack of time, some types of expressions were completely omitted from the scope of the system. Omissions include expressions that are tagged without attributes, clock times, temporal modifiers, subunits of years, and offset points of time. Examples of these types of expressions appear in (41) below.

(41) <TIMEX2>quelque temps</TIMEX2>
<TIMEX2 VAL="2002-03-21-T13:20">13h20 gmt</TIMEX2>
<TIMEX2 VAL="2000" MOD="END">fin 2000</TIMEX2>
<TIMEX2 VAL="2001-H2">second semestre</TIMEX2>
<TIMEX2 VAL="2002-SP">printemps 2002</TIMEX2>
<TIMEX2 VAL="2001">l'année dernière</TIMEX2>

While some of these omitted expressions would be easy to incorporate into the system, others would demand more complexity. Either way, the inclusion of more expressions would improve system performance, especially considering that many omitted expressions are currently being partially tagged as some other sort of recognized temporal expression. For example, those omitted expressions that include a year are currently being tagged for their year alone, lowering the system's overall score.

Evaluation

In the previous section, I discussed how well the system performs for each type of rule it supports. In this section, I will present the official scores from both French data and English data, as compared against human-annotations of each corpus.

The temporal annotation system was evaluated using a scorer that aligns tags in machine-annotated text with tags in human-annotated text. It counts the number of tags in the human-annotated text as well as the number of tags in the machine-annotated text, reporting the results in a chart under the headers "Pos" and "Act", respectively. It then calculates the number of correct tags ("Corr"), incorrect tags ("Inco"), missing tags ("Miss"), and spurious tags ("Spur"). Precision ("Prec") is calculated as the number of correctly annotated tags in the system's output divided by

the total number of tags in the machine-annotated text, and recall (“Rec”) is calculated as the number of correctly annotated tags in the system’s output divided by the total possible number of tags in the human-annotated text. The scorer reports the precision, recall, and balanced F-measure (“F”, calculated as $2PR/(P + R)$, where P is precision and R is recall) for all TIMEX2 tags, plus it breaks down results by TIMEX2 tag attributes.

When trained against the French training set, the annotation system performed fairly well on the French test set. Results from the French test set are shown in Table 7. Note that, for all results reported in this section, the F-measures for individual TIMEX2 attributes are greater than the overall TIMEX2 score. This is due to the fact that the overall TIMEX2 score accounts for attributes that the system cannot yet handle, while the individual scores reflects only those that the system should be able to annotate. For comparison purposes, I also ran the raw French training set through the tagger and scored the machine-annotated training set against the human-annotated training set. Results appear below in Table 8.

<i>Tag</i>	<i>Pos</i>	<i>Act</i>	<i>Corr</i>	<i>Inco</i>	<i>Miss</i>	<i>Spur</i>	<i>Prec</i>	<i>Rec</i>	<i>F</i>
TIMEX2	204	160	110	0	94	50	0.688	0.539	0.604
VAL	109	110	92	17	0	1	0.836	0.844	0.840

Table 7: Results on French test set

<i>Tag</i>	<i>Pos</i>	<i>Act</i>	<i>Corr</i>	<i>Inco</i>	<i>Miss</i>	<i>Spur</i>	<i>Prec</i>	<i>Rec</i>	<i>F</i>
TIMEX2	380	311	231	0	149	80	0.743	0.608	0.669
GRAN	9	9	9	0	0	0	1.000	1.000	1.000
PER	9	9	9	0	0	0	1.000	1.000	1.000
SET	9	9	9	0	0	0	1.000	1.000	1.000
VAL	222	222	170	52	0	0	0.766	0.766	0.766

Table 8: Results on French training set

Legend

Tag	Tag type
Pos	Number of tags in human-annotated data
Act	Number of tags in machine-annotated data
Corr	Correctly annotated expressions
Inco	Incorrectly annotated expressions
Miss	Expressions missed by the system
Spur	Expressions tagged only by the system
Prec	Precision = $Corr/Act$
Rec	Recall = $Corr/Pos$
F	$F_{0.5} = 2PR/(P + R)$

It is not at all surprising that the training set scored higher than the test set overall. An obvious difference between the two sets of results is that the training data contained instances of the set-denoting expressions that the learner had hard-coded into the tagging rule while none of those expressions appeared in the test set, as noted earlier. Of course, the test set is a more accurate test of the system's performance.

The system scored lower when trained on the English training set than it did with French. Results for the English test set and the English training set are shown in Tables 9 and 10, respectively.

<i>Tag</i>	<i>Pos</i>	<i>Act</i>	<i>Corr</i>	<i>Inco</i>	<i>Miss</i>	<i>Spur</i>	<i>Prec</i>	<i>Rec</i>	<i>F</i>
TIMEX2	1204	933	602	0	602	331	0.645	0.500	0.563
GRAN	59	56	46	8	5	2	0.821	0.780	0.800
PER	54	56	48	5	1	3	0.857	0.889	0.873
SET	60	56	54	0	6	2	0.964	0.900	0.931
VAL	544	546	457	81	6	8	0.837	0.840	0.839

Table 9: Results on English test set

<i>Tag</i>	<i>Pos</i>	<i>Act</i>	<i>Corr</i>	<i>Inco</i>	<i>Miss</i>	<i>Spur</i>	<i>Prec</i>	<i>Rec</i>	<i>F</i>
TIMEX2	2047	1621	969	0	1078	652	0.598	0.473	0.528
GRAN	78	76	68	1	9	7	0.895	0.872	0.883
PER	68	76	67	1	0	8	0.882	0.985	0.931
SET	78	76	69	0	9	7	0.908	0.885	0.896
VAL	890	893	742	135	13	16	0.831	0.834	0.832

Table 10: Results on English training set

There are a number of factors that may explain the difference of scores between the French and English data. I designed the system to learn rules for the types of examples available in the French training set. Tests on English data came later, though it would be wise to look to the examples in the English training set for ways to improve the system. For example, I chose to omit seasons of the year for the time being after finding that such expressions were rare (1.19% of tags) in the French corpus. In the English corpus, however, seasons of the year are slightly more frequent (2.02% of tags), so including them into the system would certainly improve its score on English data. Also, temporal expressions in the English corpus tended to be more context-dependent, which made them more difficult for the system to evaluate than were the expressions in the French corpus. In addition, the English data presented a greater need for word-sense disambiguation than did the French data. A particular problem left unsolved in the English tagger is in the rule for bare months. The learner correctly determined that the fifth month of the year is *May*, but the rule is not case sensitive when seeking month names in text. As a result, the English-trained tagger incorrectly annotates the frequently occurring verb *may* as the month of May. Changing the rule to be case sensitive would solve this problem for English, but it would lead to problems in languages like French in which month names are lower-case when occurring mid-sentence but upper-case when they begin a sentence. A better solution would be to use a part-of-speech tagger, which would disambiguate between the English verb *may* and month name *May*.

CONCLUSION

The scores reported in the previous section indicate that there remains some room for improvement in this temporal annotation system. Fortunately, the system is capable of being improved not only through handling expressions that are currently omitted but also through better processing of context-dependent expressions, like those in (35) and (36).

This learning approach to automatic temporal annotation has some weaknesses. Because the system has absolutely no knowledge of the language used in the documents it annotates, it has no access to part-of-speech information or sentence structure information that would be helpful both in determining the correct extent of a temporal expression and in disambiguating word meaning. This method is also very dependent on the quality of the training set. That is, inconsistencies or outright mistakes in the training set can hinder the learner's ability to capture the relationship between natural language text and TIMEX2 attributes, making it difficult to generate rules to explain those relationships. On the other hand, the approach also has a number of strengths. It significantly reduces the time and cost of constructing annotation systems around new languages. The evaluation results above show that the system can be easily and quickly adapted to another language given nothing more than well-annotated training data.

Future versions of this system should be able to handle the types of temporal expressions that were omitted in this version, including but not limited to clock times,

temporal modifiers, subunits of years, offset points of time, and expressions that are tagged without attributes. I would also like to change how the system deals with set-denoting expressions and improve its ability to resolve context-dependent temporal expressions. Finally, it is important to evaluate the system using corpora of additional languages (other than French and English) and genres (other than news text).

APPENDIX A: EXAMPLES OF RULES

The following are pseudocode examples of rules generated by the learner. The first is the futureRef rule, which annotates general references to some time in the future. This rule is one of the simplest rules the system learns.

```
function futureRef(untagged_string)
  altered ← False
  more_to_read ← True
  while more_to_read is True and untagged_string includes a future
  reference
    altered ← True
    insert "<TIMEX2 VAL='FUTURE_REF'>" before the match
    insert "</TIMEX2>" after the match
    tagged_string ← untagged_string with TIMEX2 inserts
    untagged_string ← string following newly inserted "</TIMEX2>"
    if untagged_string includes a future reference
      more_to_read ← True
    else
      more_to_read ← False
  end
  if altered is True
    return tagged_string
  else
    return untagged_string
```

This next rule is the dateFullDMY rule, which annotates fully specified dates containing a year, month, and date. While it is not the most complicated rule learned by the system, it does illustrate a level of complexity in word ordering and lexical meaning. To better demonstrate how the rule works, I have adapted the pseudocode specifically to French, showing regular expressions (enclosed in forward slashes) where appropriate.

```

function dateFullDMY(untagged_string)
  altered ← False
  more_to_read ← True
  while more_to_read is True and untagged_string includes
    ^\d+\w*\s\n?(janvier|février|mars|avril|mai|juin|juillet|août|septembre|oct
    obre|novembre|décembre),?\s\n?\d\d\d\d/i
    altered ← True
    temp_array ← split words in match on spaces
    rewrite the month element of temp_array to its two-digit number
    value ← year element + "-" + month element + "-" + date element
    insert "<TIMEX2 VAL="value">" before the match
    insert "</TIMEX2>" after the match
    tagged_string ← untagged_string with TIMEX2 inserts
    untagged_string ← string following newly inserted "</TIMEX2>"
    if untagged_string includes
      ^\d+\w*\s\n?(janvier|février|mars|avril|mai|juin|juillet|août|septembre
      |octobre|novembre|décembre),?\s\n?\d\d\d\d/i
      more_to_read ← True
    else
      more_to_read ← False
  end
  if altered is True
    return tagged_string
  else
    return untagged_string

```

APPENDIX B: EXAMPLES OF DOCUMENTS

The following are examples of real documents that have been annotated both by the system and by a human. The first two examples come from the French test set, and the third example is from the English test set.

Example 1

Source: Document 34

Annotator: Human

<TEXT>

JERUSALEM (AFP) - Deux personnes ont été tuées dans un attentat suicide à Jérusalem-ouest <TIMEX2 VAL="2002-03-21">jeudi</TIMEX2>, ainsi que le kamikaze, a indiqué la télévision privée israélienne.

Le nombre de blessés s'établit à 40, dont trois dans un état critique, selon la radio et la télévision publiques.

L'explosion a eu lieu devant un café-restaurant de la rue King George, tout près du carrefour avec la rue Jaffa, la plus grande artère de Jérusalem-ouest, une zone théâtre de nombreux attentats au cours des <TIMEX2 VAL="PXM">derniers mois</TIMEX2>.

La pizzeria Sbarro, cible d'un attentat suicide <TIMEX2 VAL="2001-08-09">le 9 août 2001</TIMEX2> qui avait coûté la vie à 15 Israéliens, est située à quelques dizaines de mètres de là.

Tout près se trouve également la rue piétonnière Ben Yehuda, où un double attentat suicide avait fait 11 morts <TIMEX2 VAL="2001-12-01">le 1er décembre</TIMEX2>.

Un interlocuteur anonyme a revendiqué peu après l'attentat par téléphone au nom des Brigades des martyrs d'Al-Aqsa, un groupe armé lié au Fatah, le mouvement du président palestinien Yasser Arafat.

Mohammad Hachikeh, 22 ans, originaire de Talouza, au nord de Naplouse (Cisjordanie), a commis l'attentat pour venger la mort de deux Palestiniens

tués par l'armée israélienne, a déclaré le correspondant anonyme à l'AFP.

Par ailleurs, l'armée israélienne a mené des incursions dans trois localités autonomes palestiniennes de Cisjordanie où elle a fait prisonniers une vingtaine de Palestiniens accusés "d'activités terroristes". Elle s'est ensuite retirée.

Sur le plan diplomatique, Israël et les Palestiniens comptaient poursuivre <TIMEX2 VAL="2002-03-21">jeudi</TIMEX2> leurs contacts, sous l'égide de l'émissaire américain Anthony Zinni, pour parvenir à un cessez-le-feu, après l'échec des négociations <TIMEX2 VAL="2002-03-20-TEV">mercredi soir</TIMEX2>.

Selon la radio publique israélienne, une "réunion restreinte" de la haute commission de sécurité, qui regroupe plusieurs hauts responsables des deux côtés, devrait avoir lieu <TIMEX2 VAL="2002-03-21-TEV">jeudi soir</TIMEX2>. Auparavant, un haut responsable palestinien a indiqué qu'une nouvelle réunion se tiendrait <TIMEX2 VAL="2002-03-22">vendredi</TIMEX2>, un porte-parole israélien affirmant qu'elle pourrait avoir lieu "<TIMEX2 VAL="2002-03-21-TEV">jeudi soir</TIMEX2>, sinon <TIMEX2 VAL="2002-03-22">vendredi</TIMEX2>". En outre, le Premier ministre Ariel Sharon doit s'entretenir <TIMEX2 VAL="2002-03-21">jeudi</TIMEX2> une nouvelle fois avec Anthony Zinni à Tel-Aviv, selon la radio israélienne.

En dépit de l'échec de la réunion de <TIMEX2 VAL="2002-03-20-TEV">mercredi soir</TIMEX2>, les deux parties ont fait part d'un optimisme modéré. "Nous n'avons pas réussi à parvenir à un accord durant la rencontre, mais il y a eu des discussions sérieuses sur toutes les questions de sécurité", a affirmé un haut responsable palestinien qui a requis l'anonymat. "Les voies de communications restent ouvertes, les discussions devraient reprendre <TIMEX2 VAL="2002-03-22">vendredi</TIMEX2>", a-t-il ajouté. "Nous devons encore discuter de certaines questions", a dit de son côté le porte-parole du ministère israélien de la Défense, Yarden Vatikai.

Les Palestiniens misent sur la conclusion d'un cessez-le-feu pour matérialiser une rencontre entre le président Yasser Arafat et le vice-président américain Dick Cheney, la première à ce niveau pour M.

Arafat depuis l'arrivée au pouvoir du président américain George W. Bush, en <TIMEX2 VAL="2002-01">janvier</TIMEX2>. Cette réunion pourrait avoir lieu <TIMEX2 VAL="2002-03-25">lundi</TIMEX2> au Caire, ont indiqué des sources palestinienne, égyptienne et américaine.

Toutefois, le gouvernement israélien a souligné qu'il empêcherait M. Arafat de voyager à l'étranger sauf s'il accomplit "100% d'efforts" pour la paix, comme l'a réclamé M. Cheney qui, lors d'une visite à Jérusalem <TIMEX2 VAL="2002-03-19">mardi</TIMEX2>, avait refusé de rencontrer le président palestinien. Le ministre israélien des Finances, Silvan Shalom, a indiqué à la radio publique que le cabinet déciderait <TIMEX2 VAL="2002-03-24">dimanche</TIMEX2> s'il autorise ou non Yasser Arafat à se rendre à l'étranger.

Des préparatifs pour un tel voyage sont néanmoins en cours, a indiqué un responsable palestinien. Un hélicoptère jordanien serait autorisé à se poser à Ramallah, d'où Yasser Arafat s'envolerait pour Amman, en Jordanie, probablement <TIMEX2 VAL="2002-03-24-TEV">dimanche soir</TIMEX2>, pour se rendre ensuite par avion au Caire <TIMEX2 VAL="2002-03-25-TMO">lundi matin</TIMEX2>. Arafat s'entretiendrait au Caire avec le président égyptien Hosni Moubarak, avant de se rendre en Espagne pour des discussions avec le Premier ministre José Maria Aznar, dont le pays assure la présidence tournante de l'Union européenne jusqu'en <TIMEX2 VAL="2002-06">juin</TIMEX2>.

M. Arafat irait ensuite à Beyrouth où il participerait au sommet arabe prévu <TIMEX2 VAL="2002-03-27">les 27</TIMEX2> et <TIMEX2 VAL="2002-03-28">28 mars</TIMEX2>, dont l'enjeu principal est l'examen d'une proposition saoudienne de paix pour le Proche-Orient. Après un sanglant attentat suicide palestinien dans le nord d'Israël <TIMEX2 VAL="2002-03-20">mercredi</TIMEX2> qui a tué sept Israéliens en plus de son auteur, les forces israéliennes ont été placées <TIMEX2 VAL="2002-03-21">jeudi</TIMEX2> en état d'alerte générale par crainte de nouveaux attentats, selon la radio publique.

</TEXT>

Example 1

Source: Document 34

Annotator: Computer

<TEXT>

JERUSALEM (AFP) - Deux personnes ont été tuées dans un attentat suicide à Jérusalem-ouest <TIMEX2 VAL="2002-03-21">jeudi</TIMEX2>, ainsi que le kamikaze, a indiqué la télévision privée israélienne.

Le nombre de blessés s'établit à 40, dont trois dans un état critique, selon la radio et la télévision publiques.

L'explosion a eu lieu devant un café-restaurant de la rue King George, tout près du carrefour avec la rue Jaffa, la plus grande artère de Jérusalem-ouest, une zone théâtre de nombreux attentats au cours des <TIMEX2 VAL="PXM">derniers mois</TIMEX2>.

La pizzeria Sbarro, cible d'un attentat suicide <TIMEX2 VAL="2001-08-09">le 9 août 2001</TIMEX2> qui avait coûté la vie à 15 Israéliens, est située à quelques dizaines de mètres de là.

Tout près se trouve également la rue piétonnière Ben Yehuda, où un double attentat suicide avait fait 11 morts <TIMEX2 VAL="XXXX-12-01">le 1er décembre</TIMEX2>.

Un interlocuteur anonyme a revendiqué peu après l'attentat par téléphone au nom des Brigades des martyrs d'Al-Aqsa, un groupe armé lié au Fatah, le mouvement du président palestinien Yasser Arafat.

Mohammad Hachikeh, <TIMEX2 VAL="P22Y">22 ans</TIMEX2>, originaire de Talouza, au nord de Naplouse (Cisjordanie), a commis l'attentat pour venger la mort de deux Palestiniens tués par l'armée israélienne, a déclaré le correspondant anonyme à l'AFP.

Par ailleurs, l'armée israélienne a mené des incursions dans trois localités autonomes palestiniennes de Cisjordanie où elle a fait prisonniers une vingtaine de Palestiniens accusés "d'activités terroristes". Elle s'est ensuite retirée.

Sur le plan diplomatique, Israël et les Palestiniens comptaient poursuivre

<TIMEX2 VAL="2002-03-21">jeudi</TIMEX2> leurs contacts, sous l'égide de l'émissaire américain Anthony Zinni, pour parvenir à un cessez-le-feu, après l'échec des négociations <TIMEX2 VAL="2002-03-20-TEV">mercredi soir</TIMEX2>.

Selon la radio publique israélienne, une "réunion restreinte" de la haute commission de sécurité, qui regroupe plusieurs hauts responsables des deux côtés, devrait avoir lieu <TIMEX2 VAL="2002-03-21-TEV">jeudi soir</TIMEX2>. Auparavant, un haut responsable palestinien a indiqué qu'une nouvelle réunion se tiendrait <TIMEX2 VAL="2002-03-22">vendredi</TIMEX2>, un porte-parole israélien affirmant qu'elle pourrait avoir lieu "<TIMEX2 VAL="2002-03-21-TEV">jeudi soir</TIMEX2>, sinon <TIMEX2 VAL="2002-03-22">vendredi</TIMEX2>". En outre, le Premier ministre Ariel Sharon doit s'entretenir <TIMEX2 VAL="2002-03-21">jeudi</TIMEX2> une nouvelle fois avec Anthony Zinni à Tel-Aviv, selon la radio israélienne.

En dépit de l'échec de la réunion de <TIMEX2 VAL="2002-03-20-TEV">mercredi soir</TIMEX2>, les deux parties ont fait part d'un optimisme modéré. "Nous n'avons pas réussi à parvenir à un accord durant la rencontre, mais il y a eu des discussions sérieuses sur toutes les questions de sécurité", a affirmé un haut responsable palestinien qui a requis l'anonymat. "Les voies de communications restent ouvertes, les discussions devraient reprendre <TIMEX2 VAL="2002-03-22">vendredi</TIMEX2>", a-t-il ajouté. "Nous devons encore discuter de certaines questions", a dit de son côté le porte-parole du ministère israélien de la Défense, Yarden Vatikaï.

Les Palestiniens misent sur la conclusion d'un cessez-le-feu pour matérialiser une rencontre entre le président Yasser Arafat et le vice-président américain Dick Cheney, la première à ce niveau pour M. Arafat depuis l'arrivée au pouvoir du président américain George W. Bush, en <TIMEX2 VAL="XXXX-01">janvier</TIMEX2>. Cette réunion pourrait avoir lieu <TIMEX2 VAL="2002-03-18">lundi</TIMEX2> au Caire, ont indiqué des sources palestinienne, égyptienne et américaine.

Toutefois, le gouvernement israélien a souligné qu'il empêcherait M. Arafat de voyager à l'étranger sauf s'il accomplit "100% d'efforts" pour la paix,

comme l'a réclamé M. Cheney qui, lors d'une visite à Jérusalem <TIMEX2 VAL="2002-03-19">mardi</TIMEX2>, avait refusé de rencontrer le président palestinien. Le ministre israélien des Finances, Silvan Shalom, a indiqué à la radio publique que le cabinet déciderait <TIMEX2 VAL="2002-03-24">dimanche</TIMEX2> s'il autorise ou non Yasser Arafat à se rendre à l'étranger.

Des préparatifs pour un tel voyage sont néanmoins en cours, a indiqué un responsable palestinien. Un hélicoptère jordanien serait autorisé à se poser à Ramallah, d'où Yasser Arafat s'envolerait pour Amman, en Jordanie, probablement <TIMEX2 VAL="2002-03-24-TEV">dimanche soir</TIMEX2>, pour se rendre ensuite par avion au Caire <TIMEX2 VAL="2002-03-18-TMO">lundi matin</TIMEX2>. Arafat s'entretiendrait au Caire avec le président égyptien Hosni Moubarak, avant de se rendre en Espagne pour des discussions avec le Premier ministre José Maria Aznar, dont le pays assure la présidence tournante de l'Union européenne jusqu'en <TIMEX2 VAL="XXXX-06">juin</TIMEX2>.

M. Arafat irait ensuite à Beyrouth où il participerait au sommet arabe prévu les 27 et <TIMEX2 VAL="XXXX-03-28">28 mars</TIMEX2>, dont l'enjeu principal est l'examen d'une proposition saoudienne de paix pour le Proche-Orient. Après un sanglant attentat suicide palestinien dans le nord d'Israël <TIMEX2 VAL="2002-03-20">mercredi</TIMEX2> qui a tué sept Israéliens en plus de son auteur, les forces israéliennes ont été placées <TIMEX2 VAL="2002-03-21">jeudi</TIMEX2> en état d'alerte générale par crainte de nouveaux attentats, selon la radio publique.

</TEXT>

Example 2

Source: Document 38

Annotator: Human

<TEXT>

PARIS (AFP) - Le groupe public français d'armement terrestre Giat Industries a réduit ses pertes en <TIMEX2 VAL="2001">2001</TIMEX2>, après la reprise de ses livraisons de chars Leclerc aux Emirats arabes unis, suspendues en <TIMEX2 VAL="2000">2000</TIMEX2>, mais ses perspectives restent sombres.

Giat Industries a annoncé <TIMEX2 VAL="2002-03-21">jeudi</TIMEX2> avoir enregistré une perte nette de 203 millions d'euros en <TIMEX2 VAL="2001">2001</TIMEX2> après un déficit de 282 M EUR en <TIMEX2 VAL="2000">2000</TIMEX2>. En revanche, le chiffre d'affaires a bondi de 43% à 792 M EUR, dont la moitié à l'exportation.

Cette hausse est due essentiellement à la reprise des livraisons de chars Leclerc aux Emirats arabes unis (EAU) en <TIMEX2 VAL="2001-06">juin 2001</TIMEX2>, a précisé Giat Industries dans un communiqué. Ces livraisons avaient été suspendues dans le courant de <TIMEX2 VAL="2000">l'année 2000</TIMEX2> en raison d'une divergence d'interprétation avec le client concernant le contrat signé en <TIMEX2 VAL="1993">1993</TIMEX2>.

Pour le groupe, les pertes enregistrées en <TIMEX2 VAL="2001">2001</TIMEX2> traduisent "la sous-activité du <TIMEX2 VAL="2001-H1">premier semestre de l'année 2001</TIMEX2>, avant la reprise des livraisons aux EAU mais également le surdimensionnement des ressources humaines et des moyens de l'entreprise".

Ce surdimensionnement résulte de la baisse continue du plan de charge de Giat Industries et perdure malgré les restructurations menées conformément à son Plan stratégique et social lancé en <TIMEX2 VAL="1999">1999</TIMEX2>, précise Giat.

Ce plan prévoit une réduction drastique des effectifs, à 6.500 personnes

<TIMEX2 VAL="2002" MOD="END">fin 2002</TIMEX2>, contre 10.350
<TIMEX2 VAL="1998" MOD="END">fin 1998</TIMEX2>.

Le carnet de commandes représente à <TIMEX2 VAL="2001"
MOD="END">fin 2001</TIMEX2>, <TIMEX2 VAL="P3Y6M">trois ans et
demi</TIMEX2> de chiffre
d'affaires. Ce succès n'a cependant pas dissipé l'inquiétude des syndicats
de Giat Industries, un groupe totalement dépendant des commandes militaires
et frappé de plein fouet par la réduction des budgets dans ce secteur.

L'intersyndicale a dénoncé l'"agonie" du groupe à l'issue d'un récent
Comité central d'entreprise (CCE).

La gestion du groupe public avait été épinglée en <TIMEX2 VAL="2001-
11">novembre 2001</TIMEX2> par un
rapport parlementaire sur le budget de la défense <TIMEX2
VAL="2002">2002</TIMEX2> établi par
Jean-Michel Boucheron (PS, Ille-et-Vilaine).

Ce rapport avait notamment noté la "tendance (de GIAT) à prendre des
contrats déséquilibrés à l'exportation".

Dans un rapport consacré aux industries d'armement de l'Etat publié en
<TIMEX2 VAL="2001-10">octobre 2001</TIMEX2>, la Cour des comptes
avait de son côté prôné une réduction
"conséquente" des capacités du groupe public, sans laquelle il risque selon
elle un démembrement.

En <TIMEX2 VAL="2002">onze ans depuis sa création</TIMEX2> la société
a essuyé un total de pertes
nettes supérieur à 3,66 mds EUR, dont 1,98 md EUR imputable à de mauvaises
décisions de gestion, avait souligné la Cour.
</TEXT>

Example 2

Source: Document 38

Annotator: Computer

<TEXT>

PARIS (AFP) - Le groupe public français d'armement terrestre Giat Industries a réduit ses pertes en <TIMEX2 VAL="2001">2001</TIMEX2>, après la reprise de ses livraisons de chars Leclerc aux Emirats arabes unis, suspendues en <TIMEX2 VAL="2000">2000</TIMEX2>, mais ses perspectives restent sombres.

Giat Industries a annoncé <TIMEX2 VAL="2002-03-21">jeudi</TIMEX2> avoir enregistré une perte nette de 203 millions d'euros en <TIMEX2 VAL="2001">2001</TIMEX2> après un déficit de 282 M EUR en <TIMEX2 VAL="2000">2000</TIMEX2>. En revanche, le chiffre d'affaires a bondi de 43% à 792 M EUR, dont la moitié à l'exportation.

Cette hausse est due essentiellement à la reprise des livraisons de chars Leclerc aux Emirats arabes unis (EAU) en <TIMEX2 VAL="2001-06">juin 2001</TIMEX2>, a précisé Giat Industries dans un communiqué. Ces livraisons avaient été suspendues dans le courant de l'année <TIMEX2 VAL="2000">2000</TIMEX2> en raison d'une divergence d'interprétation avec le client concernant le contrat signé en <TIMEX2 VAL="1993">1993</TIMEX2>.

Pour le groupe, les pertes enregistrées en <TIMEX2 VAL="2001">2001</TIMEX2> traduisent "la sous-activité du premier semestre de l'année <TIMEX2 VAL="2001">2001</TIMEX2>, avant la reprise des livraisons aux EAU mais également le surdimensionnement des ressources humaines et des moyens de l'entreprise".

Ce surdimensionnement résulte de la baisse continue du plan de charge de Giat Industries et perdure malgré les restructurations menées conformément à son Plan stratégique et social lancé en <TIMEX2 VAL="1999">1999</TIMEX2>, précise Giat.

Ce plan prévoit une réduction drastique des effectifs, à 6.500 personnes

fin <TIMEX2 VAL="2002">2002</TIMEX2>, contre 10.350 fin <TIMEX2 VAL="1998">1998</TIMEX2>.

Le carnet de commandes représente à fin <TIMEX2 VAL="2001">2001</TIMEX2>, <TIMEX2 VAL="PXY">trois ans</TIMEX2> et demi de chiffre d'affaires. Ce succès n'a cependant pas dissipé l'inquiétude des syndicats de Giat Industries, un groupe totalement dépendant des commandes militaires et frappé de plein fouet par la réduction des budgets dans ce secteur.

L'intersyndicale a dénoncé l'"agonie" du groupe à l'issue d'un récent Comité central d'entreprise (CCE).

La gestion du groupe public avait été épinglée en <TIMEX2 VAL="2001-11">novembre 2001</TIMEX2> par un rapport parlementaire sur le budget de la défense <TIMEX2 VAL="2002">2002</TIMEX2> établi par Jean-Michel Boucheron (PS, Ille-et-Vilaine).

Ce rapport avait notamment noté la "tendance (de GIAT) à prendre des contrats déséquilibrés à l'exportation".

Dans un rapport consacré aux industries d'armement de l'Etat publié en <TIMEX2 VAL="2001-10">octobre 2001</TIMEX2>, la Cour des comptes avait de son côté prôné une réduction "conséquente" des capacités du groupe public, sans laquelle il risque selon elle un démembrement.

En <TIMEX2 VAL="PXY">onze ans</TIMEX2> depuis sa création la société a essuyé un total de pertes nettes supérieur à 3,66 mds EUR, dont 1,98 md EUR imputable à de mauvaises décisions de gestion, avait souligné la Cour.
</TEXT>

Example 3

Source: Document W9_0046

Annotator: Human

<TEXT>

SAN FRANCISCO _ Apple Computer's acting chief executive, Steve Jobs, stunned Wall Street analysts <TIMEX2 VAL="1998-01-06-TMO">Tuesday morning</TIMEX2> by reporting that Apple Computer would report a \$45 million profit for <TIMEX2 VAL="1998-Q1">the quarter</TIMEX2> on revenue of \$1.575 billion.

The announcement, which was delivered <TIMEX2 VAL="1998-W02">a week before the computer maker is scheduled to report its first-quarter financial results for <TIMEX2 VAL="1998">1998</TIMEX2></TIMEX2>, was delivered with Jobs' characteristic flair: as a seeming afterthought that punctuated a 90-minute speech before a crowd of more than 4,000 Macintosh enthusiasts.

Analysts had expected Apple, based in Cupertino, Calif., to report a loss of 6 cents a share for the quarter, according to First Call. In <TIMEX2 VAL="1997-Q1">the year-ago first quarter</TIMEX2>, Apple reported a net loss of \$120 million, or 96 cents a share, on revenues of \$2.1 billion.

<TIMEX2 VAL="1998-01-06">Tuesday</TIMEX2>'s announcement sent the company's stock, which has traded near historic lows in <TIMEX2 SET="YES" GRANULARITY="G1W">recent weeks</TIMEX2>, soaring by 20 percent.

Apple shares closed at 18.9375, a gain of 3.0625, on volume of 16.2 million shares, the most actively traded stock in U.S. markets.

Jobs' surprise announcement was the first positive news that Apple has been able to muster since <TIMEX2 VAL="1996-Q4">the fourth quarter of 1996</TIMEX2>, when the company's <TIMEX2 VAL="PAST_REF">former</TIMEX2> chairman, Gilbert Amelio, reported a razor-thin profit of \$25 million on revenue of \$2.3 billion.

However, as even Jobs was quick to acknowledge on stage, the good news does not suggest that Apple, which has careened from crisis to crisis under a series of official and interim chief executives in <TIMEX2 VAL="P3Y">the last three years</TIMEX2>, has yet

proved that it is on a steady turnaround track.

“Who knows what's going to happen next quarter?” Jobs said. “We're going to be burning the midnight oil.”

Apple still faces skeptics who say that the company has only a slim chance of reversing the erosion of its markets by Intel-based computers that run Microsoft Corp.'s Windows operating system.

In recent quarters analysts have focused on the steady erosion of Apple's market share and have suggested that the company would continue to face a stiff challenge in persuading software developers to continue bringing out products for its Macintosh line of computers. And much of the turnaround signaled on Tuesday has been the result of a significant downsizing of the company.

“Steve Jobs has a really serious challenge going forward,” said Dan Lavin, a senior industry analyst at Dataquest, a market research company based in San Jose, Calif. “The biggest challenge is stopping the bleeding, and they showed some signs today, but they didn't prove it.”

Yet Jobs clearly made some meaningful progress on Wall Street.

Charles Wolf, a financial analyst at First Boston, said: “I did not expect this; everything was pointing to a loss. But it does show that he's got the company focused and executing.”

Jobs said Tuesday that initial sales of the company's latest computers, based on the new Powerpc G3 microprocessor produced by Motorola and IBM, suggested that the company would sell at least 1 million of the machines by the end of the year. Currently, about a third of all of Apple's sales are newer models based on the more powerful chip, he said. The company shipped 130,000 machines in this quarter _ 50,000 more than

expected, he said.

Those rates, if maintained, would suggest that Apple could sell 3 million computers in <TIMEX2 VAL="1998">fiscal 1998</TIMEX2>. Although that would represent a decline, it is above industry forecasts of about 2.6 million.

Jobs, who <TIMEX2 VAL="PRESENT_REF">now</TIMEX2> sports a beard and who walked onstage <TIMEX2 VAL="1998-01-06">Tuesday</TIMEX2> morning in a leather jacket and jeans, was a co-founder of Apple, which made the computer a widely used personal tool, and the man who created the Macintosh, which revolutionized computing by introducing a graphical interface to screens that had long been dominated by text and numbers.

He said he was feeling ``better" about his turnaround quest than he did in <TIMEX2 VAL="1997-07">July</TIMEX2> when he returned to Apple as the company's temporary chief executive.

Nibbling from a plastic bag of fruit in a hotel suite after his speech, he said he <TIMEX2 VAL="PRESENT_REF">now</TIMEX2> saw the turnaround as taking <TIMEX2 VAL="P6M">six</TIMEX2> to <TIMEX2 VAL="P9M">nine months</TIMEX2>, and he hinted that a permanent chief executive would still be found.

Separately, Fred Anderson, Apple's chief financial officer, said that the company's management team hoped that Jobs would decide to stay permanently, though he has stated publicly that he had no desire to take the job.

Jobs said he was spending the bulk of his time focusing on what he saw as the two most pressing issues facing the computer maker: corporate viability and rebuilding the company's base of software developers.

During his keynote speech, he reaffirmed his pronouncement of <TIMEX2 VAL="1997">last year</TIMEX2> that the long struggle with Microsoft was over and that the Redmond, Wash., software giant was <TIMEX2 VAL="PRESENT_REF">now</TIMEX2> a key Apple ally.

Indeed, on <TIMEX2 VAL="1998-01-06">Tuesday</TIMEX2>, Microsoft announced Macintosh Office 98, a suite of business applications that analysts believe is essential for Apple's survival in the business market.

Jobs also touted business products from Oracle Corp. and Macromedia Inc. and the newest version of the popular game Riven, which he said would be available in DVD format only on the Macintosh initially.

Apple's new relationship with the CompUSA computer retail chain was yielding promising results, he said. The chain has <TIMEX2 VAL="PRESENT_REF">now</TIMEX2> established special Macintosh sections in 57 of its stores and Jobs said that Apple's share of sales in those stores had jumped from 3 percent to 14 percent between <TIMEX2 VAL="1997-10">October</TIMEX2> and <TIMEX2 VAL="1997-12">December</TIMEX2>.

He also stressed the importance of the new version of the company's System 8 operating system, which is scheduled to ship within <TIMEX2 VAL="P1M">the next month</TIMEX2> and will include a number of new features, including the next version of Apple's Quicktime multimedia software.

Jobs elicited a chorus of boos when he said that in the new operating-system version, Microsoft's Internet Explorer would be the default World Wide Web browser.

Later, after the demonstration of Office 98, Jobs said, ``I thought I heard some of those boos turning into yahoos."
</TEXT>

Example 3

Source: Document W9_0046

Annotator: Computer

<TEXT>

SAN FRANCISCO _ Apple Computer's acting chief executive, Steve Jobs, stunned Wall Street analysts <TIMEX2 VAL="1998-01-06-TMO">Tuesday morning</TIMEX2> by reporting that Apple Computer would report a \$45 million profit for the quarter on revenue of \$1.575 billion.

The announcement, which was delivered <TIMEX2 SET="YES" GRANULARITY="G1W" PERIODICITY="F1W">a week</TIMEX2> before the computer maker is scheduled to report its first-quarter financial results for <TIMEX2 VAL="1998">1998</TIMEX2>, was delivered with Jobs' characteristic flair: as a seeming afterthought that punctuated a 90-minute speech before a crowd of more than 4,000 Macintosh enthusiasts.

Analysts had expected Apple, based in Cupertino, Calif., to report a loss of 6 cents a share for the quarter, according to First Call. In the year-ago first quarter, Apple reported a net loss of \$120 million, or 96 cents a share, on revenues of \$2.1 billion.

<TIMEX2 VAL="1998-01-06">Tuesday</TIMEX2>'s announcement sent the company's stock, which has traded near historic lows in <TIMEX2 VAL="PXW">recent weeks</TIMEX2>, soaring by 20 percent. Apple shares closed at 18.9375, a gain of 3.0625, on volume of 16.2 million shares, the most actively traded stock in U.S. markets.

Jobs' surprise announcement was the first positive news that Apple has been able to muster since the fourth quarter of <TIMEX2 VAL="1996">1996</TIMEX2>, when the company's <TIMEX2 VAL="PAST_REF">former</TIMEX2> chairman, Gilbert Amelio, reported a razor-thin profit of \$25 million on revenue of \$2.3 billion.

However, as even Jobs was quick to acknowledge on stage, the good news does not suggest that Apple, which has careened from crisis to crisis under a series of official and interim chief executives in the last <TIMEX2 VAL="PXY">three years</TIMEX2>, has yet proved that it is on a steady turnaround track.

``Who knows what's going to happen next quarter?" Jobs said. ``We're going to be burning the midnight oil."

Apple still faces skeptics who say that the company has only a slim chance of reversing the erosion of its markets by Intel-based computers that run Microsoft Corp.'s Windows operating system.

In recent quarters analysts have focused on the steady erosion of Apple's market share and have suggested that the company would continue to face a stiff challenge in persuading software developers to continue bringing out products for its Macintosh line of computers. And much of the turnaround signaled on <TIMEX2 VAL="1998-01-06">Tuesday</TIMEX2> has been the result of a significant downsizing of the company.

``Steve Jobs has a really serious challenge going forward," said Dan Lavin, a senior industry analyst at Dataquest, a market research company based in San Jose, Calif. ``The biggest challenge is stopping the bleeding, and they showed some signs <TIMEX2 VAL="PRESENT_REF">today</TIMEX2>, but they didn't prove it."

Yet Jobs clearly made some meaningful progress <TIMEX2 VAL="1998-01-06">Tuesday</TIMEX2> on Wall Street.

Charles Wolf, a financial analyst at First Boston, said: ``I did not expect this; everything was pointing to a loss. But it does show that he's got the company focused and executing."

Jobs said <TIMEX2 VAL="1998-01-06">Tuesday</TIMEX2> that initial sales of the company's latest computers, based on the new Powerpc G3 microprocessor produced by Motorola and IBM, suggested that the company would sell at least 1 million of the machines by the end of the year. <TIMEX2 VAL="PRESENT_REF">Currently</TIMEX2>, about a third of all of Apple's sales are newer models based on the more powerful chip, he said. The company shipped 130,000 machines in this quarter _ 50,000 more than expected, he said.

Those rates, if maintained, would suggest that Apple could sell 3 million computers in fiscal <TIMEX2 VAL="1998">1998</TIMEX2>. Although that would represent a decline, it is above industry forecasts of about

2.6 million.

Jobs, who <TIMEX2 VAL="PRESENT_REF">now</TIMEX2> sports a beard and who walked onstage <TIMEX2 VAL="1998-01-06-TMO">Tuesday morning</TIMEX2> in a leather jacket and jeans, was a co-founder of Apple, which made the computer a widely used personal tool, and the man who created the Macintosh, which revolutionized computing by introducing a graphical interface to screens that had long been dominated by text and numbers.

He said he was feeling ``better" about his turnaround quest than he did in <TIMEX2 VAL="XXXX-07">July</TIMEX2> when he returned to Apple as the company's temporary chief executive.

Nibbling from a plastic bag of fruit in a hotel suite after his speech, he said he <TIMEX2 VAL="PRESENT_REF">now</TIMEX2> saw the turnaround as taking six to <TIMEX2 VAL="PXM">nine months</TIMEX2>, and he hinted that a permanent chief executive would still be found.

Separately, Fred Anderson, Apple's chief financial officer, said that the company's management team hoped that Jobs would decide to stay permanently, though he has stated publicly that he had no desire to take the job.

Jobs said he was spending the bulk of his time focusing on what he saw as the two most pressing issues facing the computer maker: corporate viability and rebuilding the company's base of software developers.

During his keynote speech, he reaffirmed his pronouncement of last year that the long struggle with Microsoft was over and that the Redmond, Wash., software giant was <TIMEX2 VAL="PRESENT_REF">now</TIMEX2> a key Apple ally.

Indeed, on <TIMEX2 VAL="1998-01-06">Tuesday</TIMEX2>, Microsoft announced Macintosh Office 98, a suite of business applications that analysts believe is essential for Apple's survival in the business market.

Jobs also touted business products from Oracle Corp. and Macromedia Inc. and the newest version of the popular game Riven,

which he said would be available in DVD format only on the Macintosh initially.

Apple's new relationship with the CompUSA computer retail chain was yielding promising results, he said. The chain has <TIMEX2 VAL="PRESENT_REF">now</TIMEX2> established special Macintosh sections in 57 of its stores and Jobs said that Apple's share of sales in those stores had jumped from 3 percent to 14 percent between <TIMEX2 VAL="XXXX-10">October</TIMEX2> and <TIMEX2 VAL="XXXX-12">December</TIMEX2>.

He also stressed the importance of the new version of the company's System 8 operating system, which is scheduled to ship within the next month and will include a number of new features, including the next version of Apple's Quicktime multimedia software.

Jobs elicited a chorus of boos when he said that in the new operating-system version, Microsoft's Internet Explorer would be the default World Wide Web browser.

Later, after the demonstration of Office 98, Jobs said, ``I thought I heard some of those boos turning into yahoos."
</TEXT>

BIBLIOGRAPHY

Ferro, Lisa. "Instruction Manual for the Annotation of Temporal Expressions – Version 1.3". The MITRE Corporation: McLean, Virginia. February 22, 2001.

Ferro, Lisa, Inderjeet Mani, Beth Sundheim and George Wilson. "TIDES Temporal Annotation Guidelines Draft – Version 1.0". MITRE Technical Report MTR 00W0000094. The MITRE Corporation: McLean, Virginia. October 2000.

Jurafsky, Daniel and James H. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Prentice Hall: Upper Saddle River, New Jersey. 2000.

Mani, Inderjeet, Lisa Ferro, Beth Sundheim and George Wilson. "Guidelines for Annotating Temporal Information." *Proceedings of the Human Language Technology Conference*. San Francisco, California. 2001.

Mani, Inderjeet and George Wilson. "Robust Temporal Processing of News." *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*. New Brunswick, New Jersey. 2000 69-76.

Setzer, Andrea and Robert Gaizauskas. "A Pilot Study on Annotating Temporal Relations in Text." *Proceedings of the Workshop on Temporal and Spatial Information Processing*, Toulouse, France. 2001. 73-80.