

Light Parsing as Finite-State Filtering

Gregory Grefenstette

Rank Xerox Research Centre

6 Chemin de Maupertuis

38240 Meylan, France

Gregory.Grefenstette@Grenoble.RXRC.Xerox.com

Abstract. For a number of language processing tasks, such as information retrieval and information extraction tasks, pertinent information can be extracted from text without doing a full parse of the individual sentences. The most common restriction of the parser is to adopt a non-recursive model of the language treated, which allows an implementation of the parser using efficient finite-state tools at the cost of missing some coverage. These light parsers allow the successive introduction of symbols into the input string wherever specified regular expressions of words and/or part-of-speech tags match. Recent advances in finite-state expression compilation make writing mark-up transducers simpler, leading to quicker implementations of layered finite-state parsers. The resulting parsers are easier to create and maintain. In this article, we describe a light parsing method using recently created finite-state operators. Two applications of this parser are described: grouping adjacent syntactically-related units, and extracting non-adjacent n-ary grammatical relations. A system for evaluating the parser over a large corpus is described.

1 Introduction

For large scale text applications, a parser should be robust, rapid, and relatively accurate. In such applications as terminology extraction, lexicography, or information retrieval, one of the purposes of parsing is to recognize recurring lexical syntactic patterns and their variants. The parser used need be no more powerful than is necessary to recognize these patterns. Many such parsers [5] [7] [6] [1] [2] employ deterministic finite state recognizers over part-of-speech tagged text by first marking contiguous patterns as noun and verb groups, then marking heads within groups, and then extracting patterns between non-contiguous heads. Some of these parsers mix non-finite state procedures with finite-state recognizers, but we show here that the entire parser can be built within a finite-state framework.

Finite-state transducers can introduce markings and labels within regular patterns of tagged text. Other finite-state transducers can be composed with these marking transducers to be used as filters to extract and label a wide variety of n-ary dependency syntactic dependency relations between words. Recent advances in finite-state calculus and finite-state compilers [8] allow for more powerful and more elegant descriptions of the patterns to be recognized, making parser writing and maintenance easier.

In this paper, we briefly review finite-state transducers, and their application as text markers and filters. We then proceed in three stages: (1) using new finite-state longest-match and replacement operators, we create transducers identifying contiguous noun group and verbal group boundaries; (2) we use labelling transducers to mark the heads within each group; and (3) we create filtering transducers which label the syntactic relations between words within and across group boundaries, for example extracting verb-subject pairs. Finally, we evaluate the use of these filters over a large publically available text corpus.

2 Finite-State Transducers

Informally, a finite-state transducer¹ is a finite-state machine (visualizable as a set of nodes corresponding to states and directed arcs corresponding to transition) that consumes input and produces output with each state transition. Each arc is labeled with a symbol “under the arc” and a symbol “above the arc.” Depending upon which direction the transducer is applied in, the upper symbol can be considered as an input symbol and the lower as an output symbol, or vice-versa. If the input symbol is an epsilon, no input symbol is consumed as the arc is followed, but the output symbol is inserted in the output. The output symbol may also be an epsilon meaning that the input symbol is consumed without producing an output²

A finite-state transducer can be written as a regular expression. For example, the regular expression $xxy: a^*$ corresponds to the transducer that accepts an input of one x followed by any number of y 's and produces an output of one x and as many a 's as there were y 's in the input. This expression may also be written $xy: a^*$ since our finite-state compiler considers a single symbol without a colon as an arc with the same symbol on the input and output sides.

An arc written as $[]:x$ consumes an input x without producing an output, and an arc written as $y:[]$ outputs one y without consuming any input.

¹ See [9] for a formal treatment of the realization of deterministic transducers such as those used in this article.

² In our notation the epsilon symbol is written as a $[]$, a pair of empty brackets.

2.1 Finite-State Markers

A finite-state transducer that introduces extra symbols into an input string can be considered as a finite-state marker. We use finite-state markers to introduce brackets around nominal groups and verbal groups, and to mark head words within these groups.

2.2 Finite-State Filters

Finite-state filters are transducers which output only certain parts of the input string, setting all the other parts of the string to epsilon and possibly introducing a filter-indicating label. Once created, finite state filters, as with any finite-state transducer, can be composed or unioned.

3 Noun and verb group demarcation

Before extracting no-contiguous syntactic dependency patterns with finite state transducers, we first use marking transducers to introduce noun³ and verb group boundaries into tagged text. In his dissertation work, Debili (1982) [5] used part-of-speech precedence matrices, coupled with vectors describing possible starting and ending parts-of-speech, to describe the contour of nominal and verbal groups. Such matrices can be easily created in our finite-state calculus notation as shown in figure 1. The noun group recognizing machine described by this regular expression is given the alias `NP`. Figure 2 shows how this machine is turned into a marking transducer using the deterministic replacement [8] operator `@->` and the insertion point symbol three dots (`...`). This transducer unambiguously inserts noun group (`<NG` and `NG>`) markers around the longest instances of these groups in each sentence.

When a slightly more complete version of this group-marking transducer, as well as one for verb groups, is run over the part-of-speech tagged sentence:

```
Administration/NN of/IN 10/CD per/IN cent/NN oxy-
gen/NN to/IN the/AT ewe/NN for/IN 1/CD hour/NN
prior/NN to/IN delivery/NN did/DOD not/NOT alter/VB
the/AT surfactant/JJ properties/NNS of/IN the/AT fe-
tal/JJ tracheal/JJ fluid/NN
```

the noun and verb group markers `<NG NG>` `<VG VG>` are inserted:

```
<NG Administration/NN of/IN 10/CD per/IN cent/NN
oxygen/NN to/IN the/AT ewe/NN for/IN 1/CD hour/NN
prior/NN to/IN delivery/NN NG> <VG did/DOD not/NOT
alter/VB VG> <NG the/AT surfactant/JJ properties/NNS
of/IN the/AT fetal/JJ tracheal/JJ fluid/NN NG>
```

³ We adopt a prescriptive, rule-based, approach to such phrasal isolation, rather than a stochastic approach which recognizes small nominal chunks such as in [4]. The reasons for this are many: one need not have a hand-tagged text so that it is easier to apply the same techniques to languages other than English; we are interested in seeing how far the finite-state approach to language processing can be pushed; and by using rules, it is easier to explain and correct errors.

```
regex
[
  [ Art   => _ [ Noun ] ] &
  [ Noun  => _ [ PAdj | Prep | .#. ] ] &
  [ PAdj  => _ [ PAdj | Prep | .#. ] ] &
  [ Prep  => _ [ Art   | Noun ] ] &
]
&
[ [ Art | Noun ] [ Art | Noun | Padj | Prep ]* ]
;
END;

define NP
```

Figure 1. Finite-state compiler representation of a linked-list version of a precedence matrix, using the Arrow Rule. The double arrow operator `A => B _ C` means that whenever the regular expression `A` appears, it must be preceded by the something matching the regular expression `B` and followed by something matching the regular expression `C`. The `_` shows where the left-hand side expression `A` must appear. In the example above the regular expression `B` is always absent since only the right context is used in a precedence matrix. The line beginning with `Noun` can be read as “Noun must be followed by either a Posterior Adjective or a preposition or else end the group.” The symbol `.#.` stands for the end of the regular input string and is recognized by the finite state compiler. The line beginning `[[Art | Noun]]` intersects with preceding and states that this simplified noun group must begin with either an article or a noun, and must contain only nouns, articles, prepositions and posterior adjectives.

```
define NP
regex NP @-> <NG   ...  NG>
define MarkNGroup
```

Figure 2. Definition of the unambiguous marking transducer. The operator `A @-> B ... C` inserts the expressions `B` and `C` around the longest substrings that match the regular expression `A` in a string, in a left-to-right manner.

3.1 Marking the Heads of Phrasal Units

Once the group markers are inserted by the transducers described above into the tagged text, another transducer places head labels before certain classes of words within the groups. These additional labels indicate that the words appear in specified contexts and make the task of writing syntactic filters (see Section 4) easier and shorter.

Nominal heads are marked as being modified by prepositions (***PrepN**) or as noun phrase heads (***HeadN**) by the transducers produced by the regular expressions in Figure 3. These expressions state that the empty string [] is replaced with the symbol ***HeadN** in front of a noun or pronoun that is not followed by another noun, or by one that ends a noun group. This transducer is composed with another which replaces the ***HeadN** label with a ***PrepN** when the label is preceded by a preposition with no intervening prepositions, commas, etc.

Verbal heads are similarly marked with aspect labels inside verbal groups. Inside a verbal group, verbs will be preceded by either an active verb label ***ActV**, a passive verb label ***PasV** or as an attributive verb ***BeV**. Inside noun groups, verbs are marked as past participles ***EdV**, present participles ***IngV** or as infinitive verbs ***InfV**.

For example, the sentence:

Significant correlations were obtained between the maternal and fetal glucose levels and the maternal and fetal ffa levels.

receives, after being tagged, the following inserted labels:

```
<NC Significant/JJ *HeadN correlations/NNS NC> <VC
were/BED *PasV obtained/VBN VC> <NC between/IN
the/AT maternal/JJ and/CC fetal/JJ glucose/NN *PrepN
levels/NNS and/CC the/AT maternal/JJ and/CC fe-
tal/JJ ffa/JJ *PrepN levels/NNS NC> ./SENT
```

4 Syntactic Function Filters

The preceding section showed how finite-state transducers are used to mark contiguous group boundaries and to label head words within these groups. In this section, we show how transducers using these marks and labels can be created to extract non-contiguous syntactic n-ary dependencies. These filters take advantage of the finite-state possibility of regularly specifying what must not appear between two items.

The basic principle in describing a syntactic dependency filter between two words is to describe the structure of the syntactic context to the left of the first word, between the words, and to the right of the last word. None of this context is output (it is set to the empty string []:) by the transducing filter as the context is recognized. At the same time, a relation marking is inserted after the last word. Schematically this gives the following structure:

```
[ ]:LEFTCONTEXT
      token
      [ ]:MIDDLE
            token
            relation:[ ]
            [ ]:RIGHTCONTEXT
```

Similar syntactic filters were described and implemented in 1982 by Debili[5, p. 99]. Having introduced nominal and verbal group delimiters and labeled heads within these groups

makes it easier to write expressions that filter out syntactic functions between words. As shown in Figure 4, a simple “subject of an active verb” filter can be described as the head noun(s) in the preceding noun group. This definition will not always match the true subject, but we can estimate its accuracy as will be shown in the next section. Figure 5 shows a filter for extracting syntactic subjects of passive verb phrases.

5 Evaluation

Syntactic filters can be evaluated in the following way: (i) take any corpus, (ii) tag the corpus, (iii) mark verb and noun group boundaries, (iv) label noun and verb heads, (v) apply filters, (vi) sample results, (vii) count errors in sample. We ran these two filters over the first megabyte of AP news from 1988, and randomly chose one hundred instances of each filter output. The results are listed below. On a SPARC-20, tagging the 164,000 words took about 15 seconds of real time, inserting noun and verb phrase boundaries via non-optimized marking transducers took about 2 minutes, marking heads inside boundaries took about 3 minutes 12 seconds, and applying the union of 17 different filters took about 11 minutes. In all, this is about 10,000 words per minute. Manual evaluation showed that 7650 randomly SUBJ (Figure 4) relations from these 8000 sentences correct. 80% of the relations extracted by the Figure 5 were correct. These numbers can be improved by introducing more complicated filters, but already they provide useful indications of subcategorizations, as the frequency of passive direct objects of *killed* illustrates: people (5 times), seaman (2), villagers, vendors, teen-agers, soldiers, rebels, pilot, patient,

Here is a detailed description of the 12 errors in subject recognition that this simple SUBJ filtered produced:

- Preposition scope unrecognized — (2 of 12 errors)

Iowa [along with *Michigan*, *assessment*, and *contest*] chosen as subject of *make* in *It's evident to me that the contest in Hawaii, Michigan, and Iowa, and my own assessment of New Hampshire, make it clear that Al Haig will not be the Republican candidate in 1988*

The filter that marks prepositional dependency could not handle the paranthesization of the following sentence and *nozzle* was chosen as a subject of *came* in *Tests have resulted in "fractures ... so numerous, so large, so closely spaced and so extensive that the integrity of the entire (nozzle) came into question," said the scientists' study, obtained by the Post.*
- Subjectless imperatives unaccounted for, adverbially time expressions not implemented — (1 of 12)

day and *minute* are chosen of subjects of *listen* in *How can you conduct negotiations together when your partner runs every day and every minute to the other side and says, 'Don't listen to what Shamir says, I'll sell it to you cheaper?'*
- Tagging errors — (4 of 12)

In the version of the tagger used in these tests, any tagging errors remained unchanged and led to errors, as in the following sentence in which *charges* was tagged as a verb, with *killings* as its subject. *The killings of five Irish nationalist guerrillas and a teen-ager prompted charges that police had a shoot-to-kill policy.*

Network was tagged as a verb with *Games* as its subject below. Currently, we tokenize structurally evident proper

```

define BeginNG    "<NG"
define EndNG      "VG>"
define HeadNoun  "*HeadN"
define PrepNoun  "*PrepN"

regex
[ ] -> HeadNoun      # Add in a *HeadN label
    ||              # in the following context
    [ BeginNG |     # at the start of a noun group
      Tag ]         # or after some tagged word
    -              # Insert Label Here
    NOUN           # where it is followed by a noun
    [ ~$ NOUN ]    # which is not followed by another noun
    [INGVERB | PPART | PREP | COMMA | CC | EndNG ]
                  # before the next -ing verb, past participle,
                  # preposition, comma, conjunction or end
                  # of the noun group
;
END;
def HeadNouns

regex
HeadNoun -> PrepNoun # Replace *HeadN by a *PrepN
    ||              # label when
    [ [$ PREP] & ~$[ PREP ?* [PREP | COMMA | EndNG ]]
    ] &             # there is some preceding
                  # preposition with no intervening
                  # prep, comma or other noun group
    [ ?* Tag ]     # and put the label after a tag
    -              # Here is where old HeadNoun label is
;
END;
define PrepNouns

st clear
regex [HeadNouns .o. PrepNouns ]
fs inv

```

Figure 3. Labeling heads of nominal groups. The finite-state grammar construction $A \rightarrow B \mid C _ D$ means that A is replaced by B when it is preceded by C and followed by D . The first expression in the above figure can be read: "Insert a head noun label before any noun or pronoun that begins a noun group, or follows some other tagged word, but which is not followed by a noun." The second expression changes this tag into a prepositional label when it is preceded by a preposition without any intervening prepositions, commas, or noun phrases. These two expressions are composed using the `.o.` operator.

names such as Interactive Games Network Inc. as one token. *GAMES PEOPLE PLAY - Interactive Games Network Inc. will develop new technology that will allow viewers at home to compete against one another in live game shows and sports programs.*

The tagger was led astray by the following garden path sentence in which *struck* was considered as a finite verb with subject *supertanker*: *Shipping executives said today that a Danish supertanker struck late Thursday in an attack initially blamed on Iran now appears to have been hit by Iraqi warplanes.*

Similarly *held* is considered as a finite verb by the older version of the tagger, and *bodies* was extracted by the filter as its subject in *Jerome Woods, director the Department of Human Services, ordered the two bodies held this week, however, citing a city law that allows the government to withhold the remains if they are to be used in a "tasteless" way.*

- Non-recognition of subclauses — (4 of 12)

In the simple version of the mark-up and filters shown here, this is no recognition of subclauses, which leads to most of the incorrect subject errors as in the following sentences:

The filters extracted *Poland* as well as *Hitler* as subjects of *allow* in *Those demonstrations marked the anniversary of the 1939 Molotov-Ribbentrop pact under which Stalin agreed to allow Hitler to invade western Poland and Hitler allowed Stalin to invade eastern Poland, Lithuania, Latvia and Estonia.*

Both *efforts* and *accomplices* were extracted as subjects of *fabricate* in *Fanatics of one calling or another have in recent months tried to disrupt the relief distribution efforts and their accomplices continue to fabricate such unfounded allegations*

In the following sentence *bales* survives as subject of *incriminating* for two reasons: (i) it is not recognized as being consumed as a direct object, (ii) the conjunction of verbs to the same subject is not recognized, and (iii) *incriminating* is incorrectly seen as a continuation of the preceding verbal chain. *He said U.S. authorities had inspected 300 bales at the company's Elmhurst, Queens, warehouse but did not find incriminating evidence.*

Another error suffers from many of the same causes, as *lives* is recognized as the subject of *have* in *We just live our own lives and don't have great conflicts.*

- One other error was due to a minor coding bug in the filter.

6 Conclusion

We have presented a sequence of transducers, that composed together, provide a light parser capable of rapidly and robustly extracting common syntactic patterns from part-of-speech tagged text. By dividing the parsing task into a sequence of contiguous group marking, head labeling, and non-contiguous extracting filter transducers, we are able to write small compact rules that are easy to create and maintain. Recent finite-state calculus advances, such as the longest match operator, allow us to perform all the marking and extraction efficiently within the finite-state paradigm.

We presented an evaluation procedure for measuring the precision of relation extraction, which shows that even simple

```

#
# Subject Filter
# -----
# Input: <NG *HeadN corticosteroids/NNS NG>
#         <VG did/DOD not/NOT appear/VB to/TO *ActV
#         affect/VB VG> <NG the/AT *HeadN progress/NN
#         of/IN the/AT *PrepN disease/NN NG>
# yields:
#         corticosteroids/NNS affect/VB <SUBJ
#
regex
?:[ ] *           # Erase everything before
"*HeadN":[ ]     # Find a head noun
Token            # Keep this token
[[~ $[ "<NG" | "VG"> ]] .x. [ ] ]
                # No intervening groups
"*ActV":[ ] Token # Keep main active verb
[ ]:" <SUBJ"     # Insert <SUBJ label
?:[ ] *         # Erase everything else
;
END;
def FilterSubj

```

Figure 4. A filter for extracting subjects of active verbs.

```

#
# Passive Direct Object Filter (in Subject position)
# -----
# Ex:   <NG *HeadN amyloidosis/NN NG> <VG was/BEDZ
#         *PasV found/VBN VG>
# yields:
#         amyloidosis/NN found/VBN <PDOBJ
#
regex
?:[ ] *           # Erase everything before
"*HeadN":[ ]     # a head noun
Token            #
[[~ $[ "<NG" | "VG"> ]] .x. [ ] ]
                # No intervening groups
"*PasV":[ ]     # erase Passive label
Token            # Keep main active verb
[ ]:" <PDOBJ"   # Insert <PDOBJ label
?:[ ] *         # Erase everything else
;
END;
def FilterPassDobj

```

Figure 5. A filter for extracting passive direct objects in the subject position.

filters produce reliable results over a large corpus. Meanwhile, current research [3] in finite-state grammars is proving that an extremely high-level of accuracy is possible using finite-state rules.

REFERENCES

- [1] Steven Abney, 'Parsing by chunks', in *Principle-Based Parsing*, eds., Steven Abney Robert Berwick and Carol Tenny, Kluwer Academic Publishers, Dordrecht, (1991).
- [2] Douglas E. Appelt, Jerry R. Hobbs, John Bear, David Israel, and Mabry Tyson, 'FASTUS: A finite-state processor for information extraction from real-word text', in *Proceedings IJCAI '93*, Chambery, France, (August 1993).
- [3] J.P. Chanod and P. Tapanainen, 'A robust finite-state grammar for french', Technical report, Rank Xerox Research Centre, Meylan, France, (1996).
- [4] Kenneth Church, 'A stochastic parts program and noun phrase parser for unrestricted text', *Proceedings of the 2nd Conference on Applied Natural Language Processing*, 136-143, (1988).
- [5] Fathi Debili, *Analyse Syntaxico-Semantique Fondee sur une Acquisition Automatique de Relations Lexicales-Semantiques*, Ph.D. dissertation, University of Paris XI, France, 1982.
- [6] G. Grefenstette, 'Use of syntactic context to produce term association lists for text retrieval', in *Proceedings of SIGIR'92*, Copenhagen, Denmark, (June 21-24 1992). ACM.
- [7] Gregory Grefenstette, *Linguistic treatments applied to Information Retrieval (Traitements Linguistiques Appliquees a la Documentation Automatique)*, University of Paris XI, Orsay, France, 1983.
- [8] Lauri Karttunen, 'Directed replacement', in *Proceedings of the 34rd Annual Meeting of the ACL*, Santa Cruz, CA, (1996).
- [9] Mehryar Mohri, 'Finite-state transducers in language and speech processing', *Computational Linguistics*, **20**(1), (1994).