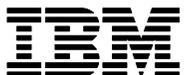


IBM Research Report

Using Slot Grammar

Michael C. McCord
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598



Using Slot Grammar

Michael C. McCord
IBM T. J. Watson Research Center

Abstract

This report describes how to use a Slot Grammar (SG) parser in applications, and provides details on an API. There is a companion report, “The Slot Grammar Lexical Formalism”. Both reports are written in a fairly self-contained way. The current one can serve as an introduction to Slot Grammar and as a kind of user guide. Topics covered include: (a) An overview of SG analysis and parse displays. (b) The use of SG in interactive mode, including an editor-based interactive mode. (c) The use of SG on whole documents, or collections of documents. (d) Descriptions and inventories of SG slots and features – both morphosyntactic features and semantic types. (e) Controlling the behavior of the parser with flag settings. (f) Tag handling and annotation methods for named entities and other text chunks. (g) The parse tree data structures. (h) The API and compilation of SG-based applications. (i) Handling user lexicons.

1 Introduction

In this report we describe how to use a Slot Grammar (SG) parser in applications. There is a companion report [13], “The Slot Grammar Lexical Formalism”, which describes SG lexicons in enough detail that the reader should be able to create such a lexicon for a new language. Both reports are written in a fairly self-contained way.

There is also a report, [12], “A Formal System for Slot Grammar”, which describes a formalism (SGF) for writing the syntax rules for a Slot Grammar. But this methodology is not currently used in SG, for reasons of efficiency.

The main current implementation of SG is in C. The SG parsers can be used in executable form, or via library (or DLL) versions, with an API described in this report. And there is a bridge of the C implementation to Java, developed by Marshall Schor in the context of UIMA (Unstructured Information Management Architecture) – see <http://www.ibm.com/research/uima>.

Currently there are Slot Grammars for English, French, Spanish, Italian, Brazilian Portuguese and German, which we call, respectively, **ESG**, **FSG**, **SSG**, **ISG**, **BPSG**, and **GSG**. More languages will come. There is a single syntactic component, `ltsyn`, for the Latin-based languages (French, Spanish,

Italian, Portuguese), with switches for the differences. We will use “XSG” to refer generically to one of the Slot Grammars.¹

The examples in this document are given mainly for English, but most of the material applies to all the languages for which we have Slot Grammars. Most of the slot and feature names are the same for all languages (are shared across languages), but we list most of the exceptions. The data structures and the API are the same for all languages.

XSG runs on a variety of platforms, including all Windows 32 platforms, AIX, Solaris, Linux, HP, and VM. Where it is relevant below, we will use terminology from Windows.

Both single-byte and Unicode versions of XSG are available. The single-byte version uses ISO 8859-1 (which is a “subset” of Unicode).

The rest of the report is organized as follows:

- Section 2, page 2: “Overview of Slot Grammar analysis structures”
- Section 3, page 6: “Running the executable”
- Section 4, page 8: “Options for parse tree display”
- Section 5, page 10: “Processing a file of sentences”
- Section 6, page 11: “The ldo interface”
- Section 7, page 11: “Features”
- Section 8, page 25: “Slots and slot options”
- Section 9, page 33: “Flags”
- Section 10, page 37: “Tag handling”
- Section 11, page 37: “Multiwords, named entities, and chunks”
- Section 12, page 47: “The data structure for SG parse trees”
- Section 13, page 52: “Punctuation and tags in SG analyses”
- Section 14, page 56: “Compiling your own XSG application”
- Section 15, page 60: “Using your own lexicons”
- Section 16, page 65: “Lexical augmentation via WordNet”

2 Overview of Slot Grammar analysis structures

As the name suggests, Slot Grammar is based on the idea of *slots*. Slots have two levels of meaning. On the one hand, slots can be viewed as names for syntactic roles of phrases in a sentence. Examples of slots are:

¹The author developed the initial versions of the Slot Grammars and continues to develop **ESG**. Claudia Gdaniec took over **GSG** and developed most of it to its current state. Esméralda Manandise took over the Latin languages syntactic component **ltsyn**, and continues to develop that, currently with special emphasis on Spanish. Contributions to the lexicons and morphology have been made by Claudia Gdaniec and Marga Taylor for German, by Consuelo Rodríguez, Esméralda Manandise, Sue Medeiros and Joaquín Salcedo for Spanish, by Esméralda Manandise for French, and by the Synthema company for Italian and Portuguese.

subj	subject
obj	direct object
iobj	indirect object
comp	predicate complement
objprep	object of preposition
ndet	NP determiner

On the other hand, certain slots (complement slots) have a semantic significance. They can be viewed as names for argument positions for predicates that represent word senses.

Let us illustrate this. Figure 1 shows an example of slots, and the phrases that *fill* them, for the sentence *Mary gave John a book*. It shows for example that *Mary* fills the **subj** (subject) slot for the verb *gave*, *John* fills the **iobj** slot, etc. One can see then that the slots represent syntactic roles.

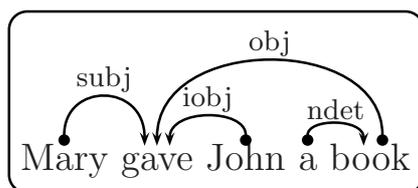


Figure 1: Slot filling for *Mary gave John a book*

To illustrate the semantic view of slots, let us understand that there is a word sense of *give* which, in logical representation, is a predicate, say $give_1$, where

- (1) $give_1(e, x, y, z)$ means “ e is an event where x gives y to z ”

The sentence in Figure 1 could be taken to have logical representation:

- (2) $\exists e \exists y (book(y) \wedge give_1(e, Mary, y, John))$

From this logical point of view the slots **subj**, **obj**, and **iobj** can be taken as names for the arguments x , y , and z respectively of $give_1$ in (1). Or, one could say, these slots represent *argument positions* for the verb sense predicate.

Slots that represent predicate arguments in this way are called *complement slots*. Such slots are associated with word senses in the Slot Grammar lexicon – in *slot frames* for the word senses. All other slots are called *adjunct slots*. An example of an adjunct slot is **ndet** in Figure 1. Adjunct slots are associated with parts of speech in the syntactic component of the Slot Grammar.

Natural language often provides more than one syntactic way to express the same logical proposition – where the variations can express extra ingredients of emphasis, topic, and the like.² Figure 2 shows an alternative syntactic way of expressing the same basic proposition as in Figure 1.

²These extra ingredients should figure in the logical representation also.

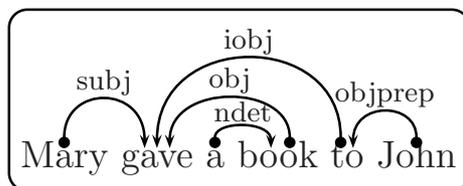


Figure 2: Slot filling for *Mary gave a book to John*

In fact, the logical representation for this sentence (not counting differences in emphasis, etc.) is the same as in (2) above. The sentence analysis in Figure 2 uses the *same* complement slot frame

- (3) (subj, obj, iobj)

as that in Figure 1. In both cases, the frame comes from the same word sense entry for *give* in the lexicon. But the *iobj* slot is filled differently in the two sentences – in the first case by the NP *John*, and in the second case by the PP *to John*. And the orderings of the slot fillers are different in the two examples. But the syntactic component of the Slot Grammar knows about these alternative syntactic ways of using slots – alternatives that lead to the same basic predication in logical form.

Because of this dual role of slots, Slot Grammar parse trees show two levels of analysis – the surface syntactic structure and the deep logical structure. The two structures are shown in the *same* parse data structure. The full form of the SG parse tree is illustrated in Figure 3, for the sentence *Mary gave a book to John*.

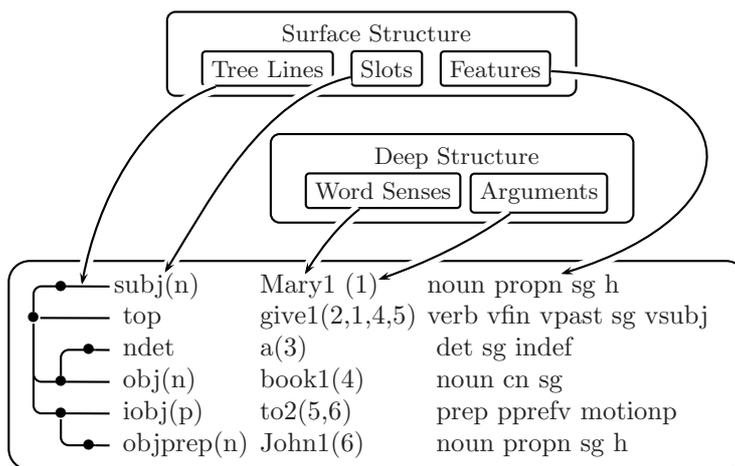


Figure 3: Ingredients of a Slot Grammar analysis structure

Note then that the surface structure of the sentence is shown in the tree lines and the slots on the left, and the features on the right. And the deep (or logical) structure is shown in the middle section through the word sense predicates and their arguments.

The lines of the parse display are in 1-1 correspondence with the (sub-)phrases, or nodes, of the parse tree. And generally each line (or tree node) corresponds to a word of the sentence. (There are exceptions to this when multiword analyses are used, and when punctuation symbols serve as conjunctions.) Slot Grammar is *dependency-oriented*, in that each node (phrase) of the parse tree has a head word, and the daughters of each node are viewed as modifiers of the head word of the node.

So on each line of the parse display, you see a head word sense in the middle section, along with its logical arguments. To the left of the word sense predication, you see the slot that the head word (or node) fills in its mother node, and then you can follow the tree line to the mother node. We describe the roster of possible slots in Section 8. To the right, you see the features of the head word (and of the phrase which it heads). The first feature is always the part of speech (POS). Further features can be morphological, syntactic, or semantic. We describe the possible morphosyntactic features in Section 7. The semantic features are more open-ended, and depend on the ontology and what is coded in the lexicon.

What are the arguments given to word sense predicates in the parse display? The first argument is just the node index, which is normally the word number of the word in the sentence. This index argument can be considered to correspond to the *event* argument *e* in (1) above (with a broad interpretation of “event”). The remaining arguments correspond to the complement slots of the word sense – or rather to the fillers of those slots. They always come in the same order as the slots in the lexical slot frame for the word sense. So for a verb, the first of these complement arguments (the verb sense’s second argument) is always the *logical* subject of the verb. Generally, all the arguments are *logical* arguments. So passivized expressions are “unwound” in this logical representation. This is illustrated in Figure 4, the parse for the passive sentence *The book was given to John by Mary*.

Note that the word sense predication `give1(4,7,2,5)` appropriately has node 7, *by Mary* (and from this *Mary*), as its logical subject, and has node 2, *the book*, as its logical object.

The type of parse display shown in this section is close to the default display for *XSG*. Actually, what we show here uses some special techniques (with `PSTricks`) in `LATEX`, especially for drawing the tree lines. This special `LATEX` form can be produced automatically by the *XSG* parser when a certain flag is turned on, and this was done for the examples of this kind in this document. The default parse display mimics the nice tree lines with ASCII characters, but it is still readable.

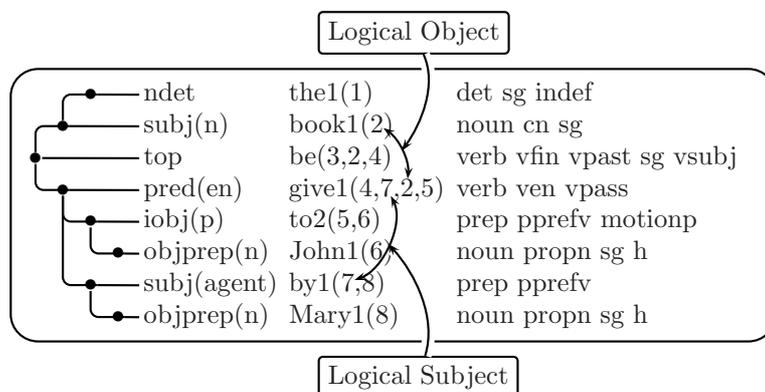


Figure 4: Parse of a passive sentence

There are several other options for parse displays (or parse output), and these are described in Section 4.

The actual C data structures for parse trees are described in Section 12.

3 Running the executable

The files needed to run *XSG* are the executable *XSG.exe* plus the lexical files. For *ESG* the lexical files are:

```
en.lxw    en.lxv
```

The lexical files for *GSG*, *SSG*, *FSG*, *ISG* and *BPSG* are similar to those for *ESG*, with *en* replaced by *de*, *es*, *fr*, *it* and *bp*, respectively.

When you run the executable *xsg* without any arguments, you will be in a loop with a prompt “Input sentence:”. We call this *interactive mode*. From this you can type in sentences (spanning several lines if you want), or other special commands. All sentences have to end with a sentence terminator, and all commands have to end with a period. Commands should be on one line. You can end the session by typing

```
stop.
```

or **Ctrl-C**. (Generally **stop** is better because it allows *XSG* to release storage, and some operating systems may not do this automatically.)

Parse tree output may be seen directly on the console, or in various editors. There are interfaces on Windows to Notepad, Kedit, and Epsilon, on Unix/AIX to Emacs, and on VM to XEDIT. The version of *XSG* that you get may be set by default to one or the other of these. If you want to direct output to the console, type:

```
-xout.
```

The command

```
+xout.
```

then causes the output to go to the file `sg.out` and then the chosen editor is invoked on that file. When you leave the file, you will be back in interactive mode.

If you have another editor, called by name *E*, then the command

```
sgeditor E.
```

will cause later output to go to that editor.

If you want to give an operating system command *C* while in the `Input sentence` loop, then you can type

```
/C.
```

For instance,

```
/dir /w.
```

would do a directory listing.

When the executable `XSG.exe` is called, it can take command-line arguments that accomplish various things. These same command-line arguments can also be passed into an initialization function `xSGstart` for the DLL form of `XSG`, as described below in Section 14.

One kind of command-line argument is `-dofile`, described below in Section 5, which allows you to run `XSG` on a text file.

Several command-line arguments allow various settings for `XSG`. The most important are *flag* settings, for the flags discussed below in Section 9. These commands are of the form:

```
-on Flag
-off Flag
```

which, respectively, turn *Flag* on or off.

The remainder of this section is a bit more esoteric, and could be skipped on first reading.

The command-line argument pair

```
-sentlen N
```

where *N* is an integer, sets the maximum length of a segment (in words, not counting punctuation or tags) that `XSG` will attempt to parse. The default is 60. It should not be set higher than 100.

The command-line argument pair

```
-timelimit N
```

where N is an integer, sets the maximum time in milliseconds that will be spent parsing a segment. The default is 15000.

One can also use command-line arguments to change default storage allocations for *XSG*. These allocation commands are of the following four forms, where N is an integer:

```
-stStorage N
-ptStorage N
-PstStorage N
-PptStorage N
```

The first, `-stStorage`, determines the number of characters in the temporary string storage buffer used by *XSG*. “Temporary” refers to the fact that the buffer is zeroed after each segment is parsed. The default value varies with the installation of *XSG*, but is in a range like 1M to 4M. The second, `-ptStorage`, determines the number of bytes in a buffer used for temporary storage of pointers and numbers. Its default value is in the range 4M to 20M. The third and fourth commands are similar, but for the main “permanent” storage areas (persisting across parses of segments). Their default values are respectively 1.5M and 2M.

The most esoteric command described here is

```
-prunedelta P
```

where P is a real number or integer. This sets an internal field `prunedelta` to P . *XSG* does parse space pruning during parsing (see the flag `prune` described in Section 9), to clear away unlikely partial analyses. The pruning is most vigorous when `prunedelta` is 0, and this is the default for most languages. When `prunedelta` is set higher, pruning is less vigorous and more parses are produced, but at the cost of efficiency (more space required and more time per parse).

4 Options for parse tree display

The default form for parse tree display has been described in Section 2. Parse trees can be displayed also in other forms. If you type

```
+deptree 0.
```

then you will see trees in a different kind of indented form, like so, for the sentence *John sees Mary*.

```

top verb vfin vpres sg vsg vsubj thatcpref
  subj(n) noun propn sg h m gname sname
    John1(1)
  see1(2,1,3)
  obj(n) noun propn sg h f gname
    Mary1(3)

```

Typing either of the following gives the default form:

```

+deptree.
+deptree 1.

```

Typing `-deptree` is equivalent to `+deptree 0`. If you type this:

```
+deptree 2.
```

then you will see trees in an XML format that may be convenient to use for some interfaces, like so:

```

<seg start="0" end="15" text="John sees Mary.">
<ph id="2" slot="top" f="verb vfin vpres sg vsg vsubj thatcpref">
  <ph id="1" slot="subj(n)" f="noun propn sg h m gname sname">
    <hd w="John" c="John" s="John1" a=""/>
  </ph>
  <hd w="sees" c="see" s="see1" a="1,3"/>
  <ph id="3" slot="obj(n)" f="noun propn sg h f gname">
    <hd w="Mary" c="Mary" s="Mary1" a=""/>
  </ph>
</ph>
</seg>

```

These are indented for easier readability. If you type

```
+deptree 3.
```

then the trees will still be in XML, but with no indentation whitespace, like so:

```

<seg start="0" end="15" text="John sees Mary.">
<ph id="2" slot="top" f="verb vfin vpres sg vsg vsubj thatcpref">
<ph id="1" slot="subj(n)" f="noun propn sg h m gname sname">
<hd w="John" c="John" s="John1" a=""/>
</ph>
<hd w="sees" c="see" s="see1" a="1,3"/>
<ph id="3" slot="obj(n)" f="noun propn sg h f gname">

```

```
<hd w="Mary" c="Mary" s="Mary1" a=""/>
</ph>
</ph>
</seg>
```

5 Processing a file of sentences

In interactive mode you can make *XSG* process a whole file by typing:

```
do InFile OutFile.
```

If you omit the *OutFile*, then output will go either to `sg.out` or to the console according as `xout` is on or off. It will send results to `sg.out` a segment at a time.

Another way to process a whole file is to invoke the *XSG* executable with arguments as follows. (The `-dofile` keyword can also be preceded by the types of command-line arguments described at the end of Section 3.)

```
XSG -dofile InFile OutFile
```

The *InFile* argument can actually be a file pattern, like `\texts*.htm`, and it will process all the files matching that pattern, sending all the output to the same output file. You can omit the *OutFile*, in which case the output will all go to `sg.out` (all the results in one step). Or you can specify *OutFile* as the empty string "", in which case the output will go to `sg.out` one segment at a time if `xout` is on, or to the console otherwise. You can also give *XSG* `-dofile` optional arguments of the following forms, appearing *after* the `-dofile` keyword.

```
-on Flag
-off Flag
-sa SubjectArea(s)
```

to turn a flag on or off or to set subject areas for the run. Available SG flags are described in Section 9. If you want to set a flag to a value *Val*, you can use an argument pair:

```
-on "Flag Val"
```

If there is more than one subject area, you should use `-sa` just once, and let its “argument” contain all the desired subject areas, enclosed in double quotes and separated by blanks, like so:

```
fsg -dofile -sa "computers instructions"
```

6 The `ldo` interface

There is an editor-based interactive interface to *XSG* which is useful in trying out variations of sentences when developing an application. It is called `ldo`. In interactive mode with *XSG* you can type:

```
ldo File.
```

This will open up *File* with the editor. Then you can place the cursor on the first line of any sentence (the cursor need not be at the beginning of the sentence) and press F6. Then *XSG* will parse that sentence and show the results in `sg.out`. When you press F3 in the editor, you will return to editing *File* in the same spot you were.

You can edit and change sentences in *File*, in order to experiment with variations of them, and again press F6 to parse them. As long as you do not save the file, when you return from `sg.out`, *File* will be in its original state (before you made a variation of the test sentence).

When you are editing *File* in `ldo` mode, you can press F3 to return to the *XSG* command line. This will leave the file in its original state, even if you have made changes – as long as you do not save the file. Of course you may want to save the file, because you have found some useful variations of sentences to work on later.

Once you have used the above `ldo` command with a given file argument, you can just type

```
ldo.
```

even after leaving *XSG* and reinvoking it, and it will return to the same file you were working on, and on the same line where the cursor was when you left the file with F3.

In preparing a test file for use with `ldo`, you should arrange it so that no two sentences share any of the same line. Start each new sentence on a new line.

The `ldo` software exists partly in the C code of the SG shell and partly in the macro language of the editor. Currently there is a editor macro program in Kedit on Windows.

7 Features

The features that appear in a Slot Grammar parse tree can be divided roughly into *grammatical* (or *morphosyntactic*) features and *semantic* features. In our formal data structure for a Slot Grammar phrase (which we describe in Section 12 below), all the features (grammatical or semantic) for each phrase are given in a single undifferentiated list of strings attached to the phrase. A phrase

always has a head word, and a feature may refer only to the head word, or to the whole phrase.

Grammatical features may represent morphological characteristics of (head) words, or may have more to do with the syntax of the whole phrase. An example of the former is **vpast**, for a past tense verb; an example of the latter is **vsubj**, which means that a verb phrase has an overt subject (in that same phrase).

Semantic features (also called *semantic types* or *concepts*) are names for sets of things in the world, and they are organized into an ontology, which contains not only the set of types, but also a specification of the subset relation between the types:

$$Type_1 \subset Type_2$$

For instance the set **h** of humans is a subset of the set **liv** of living beings: $h \subset liv$. These “things in the world” can include events, states, etc., and so for example an event referred to by a verb sense could have a semantic type. It is most typical (and useful) to mark semantic types on noun senses, but they could be marked on any part of speech.

There are two ways of dealing with semantic types in **XSG**. One is to specify them in an *ontology lexicon*, which should be named **ont.lx**. For the form of **ont.lx**, see [13]. Currently the non-English **XSG**'s use an **ont.lx**, and **ESG** can also. But the latest version of **ESG** uses a second method: The most basic types are encoded directly in **ESG** and its base lexicon; and then we use a method, described in Section 16 below, based on WordNet to augment the type marking.

For the sake of efficiency, in the current implementation the SG grammatical features, and some of the most basic semantic types, are represented internally during parsing by bit strings (or mappings onto bit positions). It is possible to do this because the grammatical features form a closed class. However, most semantic types are represented internally essentially as strings. In spite of the difference in internal (parse-time) representation of features, the SG API data structure for phrases, described in Section 12, represents both types of features as strings. This is done for the sake of simplicity of the API.

In the remainder of this section we first describe the grammatical features used in **XSG**, and then describe the most basic semantic types used for **ESG**.

We try to use as uniform a set of grammatical features across the different languages as possible. Of course some features are not applicable to all the languages. The SG system basically works with the union of all the possible features, and some of them just do not get used in all the **XSG**.

As mentioned above, the part of speech feature comes first in the list of features of a phrase. The possible parts of speech are these:

```
noun, verb, adj, adv, det, prep, subconj,
conj, qual, subinf, infto, forto, thatconj,
special, incomplete
```

The first six of these have obvious meanings. Let us go over the others:

- subconj. Subordinate conjunction, like *if, after, ...*
- conj. Coordinate conjunction, like *and, or, ...*
- qual. Qualifier, like *very, even*. Can modify adverbs and adjectives, not verbs.
- subinf. For multiwords like *in order to* and *so as to* that act like preinfinitive *to* for infinitive verbs.
- inf_{to}. The preinfinitive sense of *to*, or analogs in the other languages.
- for_{to}. The sense of *for* in *for-to* constructions like *for John to be there*.
- that_{conj}. The subordinate conjunction sense of *that* (or *que, che, daß*).
- special. A catch-all category for special “words” like apostrophe-s.
- incomplete. Used as the part-of-speech category for the top node of an incomplete parse.

Now let us look at the other possible features of a phrase. We organize these mainly by part of speech of the phrase.

Shared features. There are some features that are shared across several parts of speech. These include: **sg** (singular), **pl** (plural), **sgpl** (singular and plural), **dual** (dual), **cord** (coordinated), **wh**, and **whnom**. The last is used on determiners, nouns, and verbs; its ultimate purpose is to mark clauses, like *what you see*, that can be used essentially anywhere an *NP* can.

Verb features. We list these in alphabetical order.

- badvenadj. A past participle of such a verb does not easily fill the **nadj** slot.
Example: *said*.
- gerund. This gets marked on an *ing*-verb that has become a gerund, as in *the barking of the dogs*.
- ingprep. A verb whose *ing* form behaves like a preposition. Example: *concern*.
- invertv. A verb (like *arise* or *come*) that allows its subject on the right and a **comp PP** or *there* left modifier.
- npref. Prefers to modify nouns over verbs, as head of non-finite *VP*.
- objpref. Verb preferring *NP* object over finite clause or *that*-complement.
- q. Question clause.
- relvp. Relative clause.

- se.** German verb that takes *sein* for the present perfect.
- sta.** Stative verb.
- thatcpref.** Verb allowing both finite *VP* or *that*-complement but preferring the latter.
- transalt.** Verb (like *increase*) allowing transitivity alternation. The theme (the entity undergoing change) can be either the direct object or the subject (when no direct object is given).
- vcond.** Conditional mood.
- vdep.** Dependent clause.
- ven.** Past participle.
- vfin.** Finite verb.
- vfinf.** German infinitive verb modified by *zu*.
- vfut.** Future tense (for Latin languages).
- vimperf.** Imperfect (for Latin languages).
- vimpr.** Imperative *VP*.
- vind.** Indicative mood.
- vindep.** Independent clause.
- vinf.** Infinitive.
- ving.** Present participle.
- vlast.** *VP* whose last modifier is another *VP*.
- vobjc.** (For **GSG**) a verb with an overtly filled *obj* or *iobj*.
- vobjnom.** (For **GSG**) a verb with a nominative overtly filled *obj*.
- vpass.** Passive *ven* verb, as in *he was taken*.
- vpast.** Past tense.
- vpers1.** First person.
- vpers2.** Second person.
- vpers3.** Third person.
- vpluperf.** Pluperfect (for Latin languages).
- vpref.** Prefers to modify verbs over nouns, as head of non-finite *VP*.

- vpres. Present tense.
- vrelv. Verb that easily allows a relative clause modifier of its subject to right-extraposd.
- vsbjnc. Subjunctive.
- vsubj. *VP* with overtly filled subject.
- vthat. Relative clause with *that* as the relative pronoun.
- whatever. A clause like *whatever you see* or *whichever road you take* that is modified by an extraposd **wh-ever** *NP*.
- postcomp. Means that the verb cannot have a **comp** slot filler preceding a filler of (**obj n**).

Noun features.

- acc. Accusative.
- advnoun. A noun that can behave adverbially. Main examples: time nouns and locative nouns.
- cn. Common noun.
- cpropn. Proper noun, like “Dane” or “Italian”, that is like a common noun in denoting a class.
- dat. Dative.
- def. Definite pronoun.
- detr. Noun requiring a determiner when it (the noun) is singular.
- encprn. Can be an enclitic (for the Latin languages).
- f. Feminine.
- gen. Genitive.
- glom. A multiword noun obtained by agglomerating certain capitalized nouns in sequence.
- goodap. Can easily be a right conjunct in comma coordination even though it itself is not coordinated.
- hplmod. Noun that allows a plural **nnoun** modifier.
- indef. Indefinite pronoun.
- iobjprn. Can be non-clitic **iobj** (Latin languages).

- lmeas. Linear measure.
- m. Masculine.
- meas. Measure.
- mf. Masculine or feminine.
- mo. Month.
- nadjpn. Pronoun that can fill `nadj` and must agree with head noun (Latin languages).
- nadvn. Noun that can fill `nadv` and must agree with head noun (Latin languages).
- nom. Nominative.
- nonn. A noun that cannot have an `nnoun` modifier.
- notnnoun. A noun that cannot be an `nnoun` modifier.
- npremod. Can be a noun premodifier of a noun even though it is also an adjective.
- nt. Neuter.
- num. A number noun.
- objpprn. Can be object of preposition (Latin languages).
- objprn. Can be non-clitic obj (Latin languages).
- oreflprn. Can be use *only* as a reflexive (Latin languages).
- percent. A percent number noun.
- perspron. Personal pronoun.
- pers1. First person.
- pers2. Second person.
- pers3. Third person.
- plmod. Noun that can be an `nnoun` even when it is plural.
- posit. Position (like *middle* or *end*).
- poss. Possesive pronoun.
- procprn. Can be proclitic (Latin languages).
- pron. Pronoun.

propn. Proper noun.

quantn. Denotes a quantity (like *all* or *half*).

reflprn. Reflexive pronoun.

relnp. An *NP* that can be the relativizer of a relative clause.

tonoun. A noun, like *school*, that can by itself (without premodifiers) be the objprep of *to*, even though it is also a verb.

uif. Uninflected.

way. Manner/way.

whevern. An *NP* like *whatever* or *whichever road*.

Adjective features.

adjnoun. Adjectives like *poor* that can have a *the* premodifier and act like the head of an *NP*.

adjpass. An adjective like *delighted* that is also a past participle, but the past participle is not allowed to fill pred(en).

aqual. (For German) an adjective, like *denkbar* that can premodify an adverb (filling advpre).

compar. Comparative.

detadj. Adjectives like *next* and *last* that have an implicit definite article.

erest. Can use *-er* and *-est* for comparative and superlative. Applies to adverbs too.

lmeasadj. Adjective like *high* allowing constructions like *three feet high*.

noadv. (For German) an adjective that cannot be used as an adverb.

noattrib. Not allowed as filler of nadj.

nocompa. Not allowed as filler of comp(a).

nocompare. Not allowing comparison (for Latin languages).

nopred. Not allowed as filler of pred.

nqual. Adjectives like *medium* that allow constructions like *a medium quality car*.

post. Adjective like *available* that can easily postmodify a noun, as in *the first car available*.

soadjp. Gets put by the grammar on an adjective phrase like *so good* to allow that phrase to fill **nadv** in an *NP* like *so good a person*. This is done when the adjective (like *good*) is premodified by a qualifier marked **soqual**.

superl. Superlative.

tmadj. A time adjective like *early*.

toadj. Similar to **tonoun**.

Adverb features.

badadjmod. Preferred not to modify adjective.

compar. Comparative.

detadv. Can modify a determiner.

interj. An interjection.

initialmod. Modifies on left only as first modifier.

introadv. Adverb like *hello* that easily left-coordinates by comma-coordination.

invertadv. Adverb that allows certain constructions *Adv Verb Subj*.

loadv. Easily modifies a locational prep or adverb (particle).

locadv. Locative adverb like *above*.

noadvpre. Cannot have (qualifier) premodifier.

nopadv. Cannot modify preposition.

notadjmod. Cannot modify an adjective.

notinitialmod. Cannot appear clause-initially.

notleftmod. Cannot modify verb on left.

notnadv. Cannot premodify a noun.

notrightmod. Cannot modify verb on right.

nounadv. Can modify noun.

nperadv. Can fill **nper** slot for nouns, like *apiece* and *each*.

npost. Can postmodify a noun in slot **nadjp**.

partf. Adverb that can be a particle (also applies to prepositions that can be particles).

post. Can postmodify (like *enough*).

- ppadv. Can modify preposition (with no penalty).
- prefadv. Adverb analysis as vadv is preferred over noun analysis as obj.
- reladv. For Spanish, an adverb like *cuando* that can create a relative clause.
- superl. Superlative.
- thereprep. Adverb like *thereafter*, *thereof*,
- tmadv. Time adverb like *before*, *early*,
- vpost. Cannot modify a finite verb on the left.

Determiner features.

- all. Only for the determiner *all*.
- ingdet. Marked on possdets and *the*. Can premodify present participle verbs.
- possdet. Possessive pronoun as determiner.
- prefdet. Preferred as determiner (over other parts of speech).
- reldet. For Spanish, a determiner like *cuyo* that can create an *NP* serving as relativizer.
- the. Marked only on *the*.

Qualifier features.

- badattrib. If it modifies an adjective, then that adjective cannot fill nadj slot.
- c. Modifies only comparative adverbs.
- post. Can postmodify.
- pre. Can premodify (the default).
- soqual. See soadjp for adjectives above.

Subordinate conjunction (subconj) features.

- assc. For *as* (or analogs in other languages).
- comparsc. For *assc* or *thansc*.
- finsc. Allows only finite clause complements (filling scomp).
- notleftsc. A clause with this as head cannot left-modify a clause (as vsubconj).
Example: *for*.

okadjsc. Allows adjective complement.

oknounsc. Allows noun complement.

oknsubconj. Can fill nsubconj.

poorsubconj. Preferred not as subconj.

sbjncsc. For the Latin languages: Suppose a `subjconj` *S* has a finite clause complement *C* (*C* is filler of `scomp`). Then *S* must be marked `bjncsc` if *C* is marked `vsbjnc` (subjunctive) but not `vind` (indicative). And if *S* is marked `sbjncsc` and *C* is marked `vsbjc`, then remove `vind` from *C* if it is present.

thansc. For *than* (or analogs in other languages).

tosc. Allows `inf`to complement.

whsc. A wh-subconj (like *whether*).

Preposition features.

accobj. (For German) allows the `objprep` to be accusative.

adjobj. (For German) allows the `objprep` to be an adjective or past participle phrase.

asprep. For *as* and analogs in other languages.

badobjping. Cannot have a `ving objprep` (under certain conditions).

daprep. For German. For *PPs* like *dabei* and *darauf*. A word *da+Prep* is unfolded to a *PP* with head *Prep* marked `daprep` and `objprep` filled by *es*.

datobj. (For German) allows the `objprep` to be dative.

genobj. (For German) allows the `objprep` to be genitive.

hasadjobj. (For German) the `objprep` is an adjective or past participle phrase.

infobj. (For Latin languages) has an `objprep` that is an infinitive *VP*.

locmp. Used for multiword prepositions that fill `comp(1o)`.

motionp. Prefers not to fill `comp` if the matrix verb is marked `sta`.

nonlocp. Cannot be a filler of `comp(1o)`.

notwhpp. Cannot have wh `objprep` (under certain conditions).

- pobjp.** The **objprep** can be a *PP* itself. Example: *from*.
- ppost.** The preposition can follow the **objprep**. If the preposition is marked **ppost** but not **ppre**, then the preposition *must* be on the right.
- ppre.** The preposition can precede the **objprep**. There is no need to mark this feature on the preposition, in order to allow the preposition to be on the left, unless the preposition is marked **ppost**.
- pprefn.** Prefers to modify nouns.
- pprefv.** Prefers to modify verbs.
- preflprn.** Marked by the grammar on a *PP* when the **objprep** is a reflexive pronoun.
- relpp.** A *PP* that can serve as a relativizer for a relative clause.
- staticp.** Preposition like *in* or *on* that is used normally to represent static location vs. goal of motion (as with *into* and *onto*).
- timep.** Common modifier of time nouns, as in *two days after the meeting*.
- timepp.** For certain *PPs* where the **objprep** is a time noun.
- woprep.** Like **daprep**, but with *wo* instead of *da*.

Basic semantic types.

Here we list and briefly describe the most basic semantic types used with **ESG**. They are listed in alphabetical order. Most of the types are marked on nouns, and just a few (named specifically) are marked on verbs.

- abst.** Abstraction.
- act.** Act by a human.
- ameas.** Unit of area measure.
- anml.** Animal (not including humans).
- artcomp.** Artistic composition – literary, musical, visual, dramatic, etc.
- artf.** Artifact.
- century.** Named century like “300 B.C.”
- chng.** Event of change.
- cmeas.** Unit of volume (cubic) measure.
- cognsa.** Cognitive state or activity.

coll. Collection.

collectn. Collective noun dealing overtly with a set – like “set”, “group”, “family”. Subtype of **coll**.

cpropn. Meant to suggest “common proper noun”. A **propn**, like “German” that can name a class, and so behaves also like a common noun.

cst. State or province in a country.

ctitle. A postposed company title, like “Inc.” or “Co.”.

ctry. Named country (**propn**).

cty. Named city.

discipline. Branch of knowledge.

doc. Document.

dy. Day. Named weekdays like Tuesday or named holiday days.

emeas. Unit of electromagnetic measure.

evnt. Event.

f. Female.

feeling. Feeling, emotion.

geoarea. Geographical area.

geoform. Geological formation.

geopol. Named geopolitical unit – like **ctry**, **cst**, **cty**.

gname. Human given name.

h. Human individual.

hg. Human group.

imeas. Unit of illumination measure.

inst. Instrument – artifact used for some end.

langunit. Language unit, like word, discourse, etc.

liv. Living thing.

lmeas. Unit of linear measure..

loc. Pronoun for location, like “here”, “anywhere” or certain general locational common nouns like “outside”.

- location. Point or region in space.
- locn. Named location – like a geopol or an ocean.
- locnoun. Noun like “here” or “east” that can function adverbially or fill comp(lo).
- m. Male.
- massn. Mass noun.
- meas. Unit of measure.
- mmeas. Unit of money.
- mo. Named month, like “October”.
- name. Usual sense.
- natent. Natural entity (not made by humans).
- natlang. Named natural language.
- natphenom. Natural phenomenon.
- ntitle. A postposed part of a human name like “Jr.” or “Sr.”.
- org. A human organization. Subtype of hg.
- percent. Percent expression.
- physobj. Physical object.
- physphenom. Physical phenomenon.
- process. Sequence of events of change.
- professional. Person with a profession requiring higher education.
- propcn. A common noun, like “society” or “mathematics”, that can name a single entity and function like a propn – sometimes capitalized, especially in older writing.
- property. (Synonym of “attribute”).
- ptitle. A postposed human title, like “Ph.D.” or “M.D.”
- quantn. Quantity pronoun like “more”, “half”, or ordinal number, or collectn noun.
- relig. A named religion.
- r1ent. Role entity. An entity (usually person) viewed as having a particular role – includes many human roles like leader, artist, engineer, father,

rmeas. Relation between measures, like rate or scale.

saying. Expression or locution.

sayv. Verb of saying.

subst. Substance.

smeas. Unit of sound volume measure.

sname. Human surname.

socialevent. Human social event.

speechact. Speech act, e.g. request, command, promise.

strct. Structure (thing constructed).

title. Human title, like “President” or “Professor”.

tm. Time noun, like “year” or “yesterday”.

tma. Time noun, usually a pron or propn, or behaving as such – e.g. “now”, “tomorrow”.

tmdetr. Time noun, like “season”, that requires a determiner in order to act adverbially.

tmeas. Unit of temperature measure.

tmperiod. Time period.

tmrel. Time noun allowing certain finite clauses as **nrel** modifiers.

trait. Distinguishing property (usually applied to persons).

transaction. Act in conducting business.

ust. U.S. state. | cst

vchng. Change (verb type).

vchngmag. Change magnitude/size (verb type).

vcreate. Cause sth to exist or become (verb type).

vmove. Change locations (verb type).

wlocn. Named water mass, like “Atlantic Ocean”.

wmeas. Unit of weight measure.

yr. Year. Specific year, usually given numerically.

8 Slots and slot options

The SG parse tree shows for each phrase (or tree node) the slot that is filled by the phrase. Slots are of two kinds – *complement* slots and *adjunct* slots. As indicated in Section 4, the complement slots correspond to logical arguments of the head word sense of the phrase; they are specified in the lexical entry for the word sense. The possible complement slots for verbs are these six:

subj	<i>subject</i>
obj	<i>direct object</i>
iobj	<i>indirect object</i>
pred	<i>predicate complement</i>
auxcomp	<i>auxiliary complement</i>
comp	<i>complement</i>

We discuss these in more detail below.

The fillers of adjunct slots are like “outer modifiers”. In logical form, they often predicate on the logical form of the phrase that they modify. For instance the determiner slot **ndet** for an *NP* may contain a quantifier that is like a higher-order predicate on the rest of the *NP*. Adjunct slots are associated with parts of speech, and are specified in the syntactic component of *XSG*.

An SG phrase also shows the *slot option* chosen for a given slot. Slot options correspond roughly to the phrase category (part of speech) chosen for the filler of the slot. A given slot may be filled by phrases of various categories. For instance the *iobj* slot in English can normally be filled by either an *NP* or a *to-PP*, as in these two examples:

Mary gave John the book.
Mary gave the book to John.

The lexical entry for a word sense shows, for each of its complement slots, what slot options are allowed for that slot. The semantic idea is that (as mentioned) the slot corresponds to a logical argument, but the list of options for the slot indicates how the argument can be realized syntactically. (The slot itself also carries syntactic information though.) The current possible SG slot options are these:

a, agent, aj, av, bfin, binf, dt, en, ena,
 fin, fina, finq, finv, ft, ger, gn, inf, ing, io, it,
 itinf, itthatc, itwh, lo, n, na, nen, nmeas, nop, nummeas,
 p, padj, pinf, pinfd, prflx, prop, pt, pthatc, pwh,
 qt, rflx, sc, so, thatc, v, wh

We will not describe these here because usually (though not always) the feature list of the phrase contains the information provided by the slot option – especially since the feature list includes the part of speech of the phrase. But see

[13] for a description of the main options. Generally, one should think of slot options as providing control for parsing itself. There is one case at least where it is important to look at the slot option: The `comp` slot for verbs can have an `n` option, fillable by *NPs*, for object complements like this:

They elected $\underbrace{\textit{Ellen}}_{\text{obj(n)}}$ $\underbrace{\textit{president of the company.}}_{\text{comp(n)}}$

But `comp` also allows an option `io`, fillable by *NPs*, for an alternative form of the indirect object, like this:

They took $\underbrace{\textit{John}}_{\text{comp(io)}}$ $\underbrace{\textit{the contract.}}_{\text{obj(n)}}$

Since in both cases `comp` is filled by an *NP* – even a human *NP* – it is worth looking at the distinction of options `n` vs. `io`.

Now let us look at the *SG* slots. We organize the description by part of speech. It is worth noting that most slots – especially complement slots – allow several options. In some of the description, we will be a bit imprecise by not distinguishing between a slot and the filler of the slot. For instance, in a statement like “The `subj` agrees with the finite verb”, we really mean: “The filler of `subj` agrees with the finite verb”.

Verb complement slots. We will elucidate these in part in terms of the thematic roles *THEME*, *LOCATION*, *GOAL*, and *AGENT*, in the sense of Gruber [3] and Jackendoff [4]. In verbs that describe a change, the *THEME* is what changes, and it changes to the *GOAL* state. The *AGENT* is the participant that carries out the change. In verbs that describe a state, the *THEME* is what the state is predicated of and the *LOCATION* is that state.

`subj`. Most readers of this document will be familiar with the subject slot, but let us make a few comments. If a verb has an *AGENT*, then this will normally fill `subj`. The subject may also be *THEME*. The default position of the subject in our European languages is before the verb in an active declarative clause. The subject is filled overtly (in these languages) only in finite clauses, and then it must agree in person and number with the finite verb. In non-finite *VPs*, we often show the logical subject in the word sense predication. The `subj` slot may be filled by several different categories of phrases besides noun phrases, as in:

Seeing him was difficult.
That he did that is amazing.
Whether he can go is an open question.

obj. If a verb has a THEME and a filled obj, then that obj is normally the THEME. The obj may be filled by several different categories of phrases, as in:

He believes John's story.
He believes in John.
He believes that John was there.
He believes what John said.
John said, "I was there".
He likes seeing John.
He believes John to be honest.

iobj. The indirect object normally has a GOAL role. In English, it can most typically be filled by both an *NP* and a *to-PP*. But some verbs, like *ensure*, can take only an *NP* iobj. And others, like *confess*, take only a *to-PP* iobj. The iobj sometimes allows other prepositions, as in:

She bought some books for me.
She bought me some books.

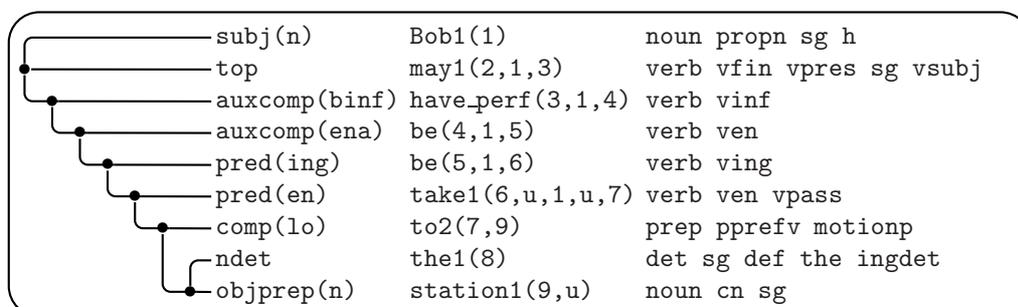
An *NP* iobj normally takes the dative case in our European languages.

comp. In the thematic roles, the basic idea of comp is that it is the GOAL in verbs of change and the LOCATION in verbs that describe a state. The iobj slot could actually be folded into the comp slot, and in fact the option io for comp is more or less an alternative way of expressing iobj(n), as mentioned above. But we use iobj, partly for the sake of familiarity. As with the other verb complement slots, comp allows several options. We gave *NP* examples above. The most typical comp is a *PP*. Examples:

Alice drove Betty to the store.
 obj(n) comp(lo)
Alice drove Betty to distraction.
 obj(n) comp(lo)
Alice drove Betty crazy .
 obj(n) comp(a)

pred. In English, pred is a slot only for the verb *be*. It can be filled by a very wide range of categories of phrases, illustrated in part by the following:

Bob is a teacher.
Bob is happy.
Bob is in love.
Bob is to leave tomorrow.
Bob is leaving tomorrow.
Bob was taken to the station.

Figure 5: *Bob may have been being taken to the station.*

In our other European languages, **pred** is used for some more verbs besides the analogs of *be*; for instance for German it is used for *werden* as well as *sein*.

auxcomp. This slot is used in English mainly for the modal verbs, where **auxcomp** is filled by a bare infinitive, for the auxiliary *do*, which also takes a bare infinitive, and for the perfect sense of *have*, where the filler is an active past participle. Two of these cases are illustrated in Figure 5, which also contains the use of **pred** for the progressive and the passive.

Verb adjunct slots. The names of most of these slots begin with “v”. For most of these, instead of giving a general definition, we will just give one or more examples.

vadv. *He always likes chocolate.*

vprep. *In that case, please buy some chocolate.*

vsubconj. *If the shoe fits, wear it.*

vnfvp. The nfv stands for “non-finite verb phrase”.

Using the mouse, select the chocolate icon.

Made in Switzerland, this chocolate is sold in many countries.

To select the chocolate icon, use the mouse.

vfvp. The fvp stands for “finite verb phrase”.

However you do it, select the chocolate icon.

Had you selected this kind of chocolate, you would be happier.

vforto. *For John to be there on time, we’ll have to call.*

vcomment. *This chocolate, I think, is best.*

vadjp. *Happy with the results, she left.*

- vvoc. John, can you hear me?
 vadj. happy-seeming
 vdet. his seeing the car
 vnoun. user-defined
 vrel. Someone appeared who I didn't expect.
 vextra. It was assumed that she would arrive Tuesday.
 vpreinf. zu sehen
 vnp. Votre ami, aime-t-il le chocolat?
 vinf. Trouver ce chocolat, c'est très difficile.
 vdat. Mi mangio una mela.
 vnthatc. Er hat keine Erkenntnis gehabt, daß sie da war.
 whadv. When did you eat?
 whprep. At which table did you sit?

Noun complement slots.

There are four of these. The most important are:

nsubj, nobj, ncomp

For deverbal nouns, these slots correspond to the verb slots:

subj, obj, comp

– allowing for the view that *iobj* is a kind of special case of *comp*. These noun slots have roughly the same range of possible slot options as the corresponding verb slots do. But the *n* option for the noun slots is filled by an *of*-PP instead of an NP.

For example we have the correspondence:

<u>John</u>	knows	<u>that Bill was there</u>
subj(n)		obj(thatc)
<u>John's</u>	knowledge	<u>that Bill was there</u>
nsubj(n)		nobj(thatc)

John, my brother
Paris, the capital of France

- nloc. *Paris, France*
 nprep. *the man in the house*
 ngen. *ein Freund meines Vaters*
 nrel. *the man that I saw yesterday*
 nrela. *the man I saw yesterday*
 nnfvp. *anyone seeing John*
 anyone seen by John
 the first person to see John
 nsubconj. *the situation as described by John*
 ncompar. *more money than he has*

Adjective complement slots. There is only one – aobj. Examples:

happy with the result
happy to see the result
happy that the result was good

Adjective adjunct slots.

- adjpre. *very happy*
 anoun. *user-friendly*
 adjpost. *happy enough*
 acompar. *happier than John*
 aprep. *red in the face*
 asothat. *so red that he glowed*
 arel. *der Rote, der auf dem Tisch steht*

Adverb complement slots. There is only one – avobj. Examples:

enough for me
enough to accomplish the result
enough that the result was good

Adverb adjunct slots.

advpre. very quickly

advpost. quickly enough

advinf. how to do it

avsothat. so quickly that it smoked

avcompar. more happily than John

Preposition complement slot: objprep. Examples:

in the house

in seeing the house

in what you see

Preposition adjunct slots.

padv. out in the yard
expressly for John

pvapp. daran, daß er geht

Determiner adjunct slot: dadv. Example: almost all people

Complement slot for subconj: scomp. Example: if the shoe fits

Adjunct slot for subconj: scadv. Example: only while she eats

Complement slot for infcto: tocomp. Example: to eat chocolate

Adjunct slot for infcto: toadv. Example: only to eat

Complement slots for forcto.

forsubj. for John to write the book

forcomp. for John to write the book

Complement slot for thatconj: thatcomp. Example: that he wants to go

Complement slot for subinf: subinfcomp. Example: in order to see the show

9 Flags

The SG system has several *flags* that can be set to control various aspects of parsing. Each flag can have an integer value, which is normally either 1 (flag is on) or 0 (flag is off), but a couple of flags can have other integer values. The flag `deptree`, discussed above, can have values 0, 1, 2, or 3. Default values for all flags are set by `XSG` at initialization time.

To set a flag *F* on (or off) in interactive mode, you can type `+F` or `-F` respectively. Or you can type `+F n` to set the value to any non-negative integer value *n*. (So `-F` is equivalent to `+F 0`.)

You can also set flags when you compile your own `XSG` application, by using the API function `xSGsetflag`, described on page 60 in Section 14.

Here is a list (in alphabetical order) of the most useful flags and their meanings. We will usually explain what the effect of having the flag *on* is.

- `all`. Process all of the top-level parses. When `all` is off, only the first (best-ranked) parse is processed. *Processing* a parse means displaying it, if the flag `syn` is on, and calling any post-parse hook functions.
- `allcapprop`. The general idea of `allcapprop` is that a word that consists of all capital letters will be treated as a proper noun. But more precisely it is this: Suppose (1) the flag is on, (2) *W* is a word consisting of all capital letters and having at least two letters, and (3) the containing segment does not consist totally of capital letters. Then the system will use only *W* itself in morphology and lexical look-up; it will *not* do morphology and look-up on the lower-case form of *W*. The latter action is the default. For instance, if the input segment contains *MAN*, the system will look up both “MAN” and “man”. If `allcapprop` is on (and the segment is not all caps), then it will look up only “MAN”. This flag is off by default.
- `allcapprop1`. This has the following effect. Suppose (1) the flag is on, (2) *W* is a word consisting of all capital letters and having at least two letters, (3) the containing segment does not consist totally of capital letters, and (4) there is no lexical analysis of *W* itself. Then the system will create a proper noun analysis for *W*. And, if `allcapprop` is off, the system will add any lexical analyses it can obtain for the decapitalized form of *W*. This flag is off by default.
- `capng`. This causes capitalized noun groups to be agglomerated into multi-word proper nouns. The noun group can contain (capitalized) common nouns, proper nouns, or adjectives. The agglomeration will not take place unless the segment is an *lcseg*; this means that the segment contains some “content words” (nouns or verbs) in lower case. The flag `capng` is off by default.
- `capprop`. This is similar to `capng` but does not produce as many proper noun agglomerations. If the segment is an *lcseg* then it agglomerates

sequences of capitalized words that have some proper noun analysis. Else it agglomerates only capitalized words that have a *unique* analysis and this analysis is a proper noun. This flag is on by default.

captprop. This is the strongest of schemes for turning capitalized words into proper nouns. The basic idea is that every capitalized word *W* becomes a proper noun, but there are several exceptions, as follows: (1) The segment is not an *lcseg*; (2) *W* is the first word of the sentence, or the first word after a quote; (3) *W* is a pronoun or already has a proper noun analysis; (4) *W* has one of the semantic types:

```
tma, propcn, meas, title, hlanguage, st_people,
st_religion, st_discipline, mmeas, wmeas, emeas, cmeas
```

This flag is off by default.

colonsep. Basically this causes colons to be treated as segment terminators. It is on by default. But if in addition the flag **linemode** is on, then the flag **lncolonsep** must be off in order for colons to be treated as terminators. When colons are not terminators, they will be treated as punctuation conjunctions.

deptree. Controls type of parse tree display. Discussed above.

displinecut. If the current input is within a *display tag* (see **dodisplays**), then newlines break segments, i.e., segments cannot span across lines. On by default.

dodisplays. Certain tags are designated *display tags* in the shell (and which ones they are depends on the formatting language). The flag **dodisplays**, which is on by default, enables processing (parsing, etc.) of text within display (begin and end) tags.

doshowstat. Show parsing statistics for a file. Output goes at the end of the output file. On by default.

echoseg. Echo (print out) the input segment for each parse. On by default.

echosegext. Echo segment-external material in output. Off by default.

forcenp. This is used in **ESG**. When it is on, only the NP analyses of a segment will be produced, as long as there exists at least one NP analysis. Off by default.

fullfeas. Activates display of additional phrase features in parse output, namely the features:

```
le1, le2, le3, le4, ri1, ri2, ri3,
xtra, cord, comcord, unitph, lcase, glom
```

Among the `len`, it shows only the strongest one – the one with greatest n which the phrase has as a feature. Similarly for the `rin`. This flag is off by default.

`html`. Enables processing of HTML documents. On by default. The flag `sgml` should also be set on when `html` is on.

`ibmidoc`. Enables processing of IBMIdDOC documents (used in some IBM manuals). Off by default. The flag `sgml` should also be set on when `ibmidoc` is on.

`limitall`. Suppose the value of this flag is the integer N . If the flag `all` is on, then only the first N parses of the current segment will be processed (if $N \leq 0$, this means that no parses will be processed). If `all` is off, then exactly the first parse will be processed, no matter what `limitall` is set to. The default value of `limitall` is a “large integer” (like 1000000).

`linemode`. Causes newlines to break segments (so that segments cannot span across lines). Useful for regression testing. A segment (on a line) need not have a segment terminator (like a period, question mark, etc.). Off by default.

`lncolonsep`. It is on by default. See flag `colonsep` for the effect of this flag.

`noparse`. If this flag is turned on, then only the morphological part of parsing will be done; the chart parsing will be skipped.

`otext`. Enables processing of Otext documents. The flag `sgml` should also be turned on when `otext` is on. Off by default.

`predargs`. This is on by default, and in parse displays, it enables word senses to be shown with their arguments, like:

```
believe2(2,1,3,4)
```

When it is off, you would see just `believe2` in the head sense position.

`predargslots`. When this is turned on, it causes word sense predications to be displayed with slot names attached to the arguments, like so:

```
believe2(2,subj:1,obj:3,comp:4)
```

`prefnp`. This is used in **ESG**. When it is on, short top-level segments that have *some NP* analysis are preferred as *NPs* – the shorter, the more preference. This is implemented in `topeval` by decrementing the `eval` field by an amount that increases with shorter segments. This flag is off by default. It can also be turned on with a value different from 1 (or 0). When it is given a non-zero value, that value will be used as a multiplier in a certain formula that rewards NP analyses. So for example a value of 3 will produce 3 times the reward of a value of 1.

- printsentno.** When this is on (and it is on by default), and a file is parsed, the segment number of each segment will be printed before the segment string in the output file.
- prune.** When this is on (and it is on by default), the parsing algorithm prunes away (discards) partial analyses whose parse scores become too bad. We will not describe the details of parse scoring in this report, but see [12]. If you turn this flag off, the parser generally produces many more parses. If not enough space is allocated for *XSG* and **prune** is off, then there may be problems for parsing longer sentences. An alternative to turning **prune** off is to increase the *XSG* field **prunedelta**, as described in Section 3.
- semicolonsep.** Causes semicolons to be treated as segment terminators. It is on by default. When it is off, semicolons will be treated as punctuation conjunctions.
- sgml.** Enables processing of SGML texts. It is on by default. Even when it is on, plain text is handled because plain text is just viewed as SGML text with no tags.
- showaopts.** This is off by default, but when it is turned on, the chosen option *o* for each *adjunct* slot *s* will be shown in parses in the form *s(o)*. So for example for the determiner **the**, one would see the slot **ndet(dt)** instead of just **ndet**.
- shownumparses.** In parse output, shows the total number of parses obtained. On by default.
- shownumsent.** When this is on and a file is parsed, the segment number of each segment will be printed to the console. On by default.
- showopts.** When this is turned on, the chosen option *o* for each complement slot *s* will be shown in parses in the form *s(o)*. On by default.
- showsense.** In word sense predications in parse displays, this causes sense names of the words to be used as the predicates. It is overridden by **showssense**. If both **showsense** and **showssense** are off, citation forms will be used for the predicates. On by default.
- showslots.** Show the available (unfilled complement) slots of each phrase in the parse tree. Off by default.
- showssense.** For the predicate of each word sense predication in parse displays, this shows a deeper kind of sense expression for some word senses (details not given here). Off by default.
- spacelinecut.** If this on, then each line of input text that consists only whitespace will break the current segment. Off by default.
- sgsyn.** Causes parse trees to be printed out. On by default.

- `timit`. Causes tracing of time used for analyzing each sentence. On by default.
- `toktrace`. Causes display of the tokens for a segment that are produced by the tokenizer. See Section 13 for the description of the `token` data type.
- `xout`. When it is on, results go, in interactive mode, to the file `sg.out`. When it is off, they go to the console. Off by default, except on VM.

This completes the inventory of flags.

10 Tag handling

XSG can read texts in various formats. One of these of course is plain text (without any tagging). In addition, **XSG** can handle various forms of SGML or XML and the IBM BookMaster document format. The main two forms of SGML known to **XSG** are HTML and IBMIdDoc. See the flags in Section 9 for switching to these tag systems. Adapting to a new form of SGML is a matter of specifying for **XSG** certain classes of tags for that DTD (this currently cannot be done directly by the user).

Generally **XSG** does better with tagged text than with plain text, for at least three reasons:

(1) Sometimes in plain text it is not clear to **XSG** where segments end, because segments like section titles are usually written without end punctuation. Tags will usually show **XSG** where such segments end.

(2) In tagged text, tag pairs within segments like those that indicate font changes, italics and boldface are often useful to **XSG** in delineating subphrases of segments.

(3) Some tags can provide clues to **XSG** as to the kind of segment that follows the tag. For example, headings or titles are often noun phrases as opposed to sentences, and **XSG** can pay attention to heading tags like the HTML `<H1>` and parse the following segment with a preference (though not strict requirement) for an NP analysis.

(4) Sometimes parts of documents are not meant to be processed by the parser or NLP system. For example an MT system should not try to translate a Javascript section of an HTML document. In tagged text, such sections are usually delimited by certain tags, which **XSG** can pay attention to.

11 Multiwords, named entities, and chunks

XSG has several strategies for handling multiwords, named entities, and special kinds of tokenization. We discuss these in this section.

11.1 Normal lexical multiwords

The *XSG* lexicons contain many multiword entries, which can be of any part of speech. Multiwords always have a *headword*, which can be inflected. The headword is by default the last word in the multiword, but this can be overridden by putting an = sign in front of the desired headword, as in this entry for “attorney general”:

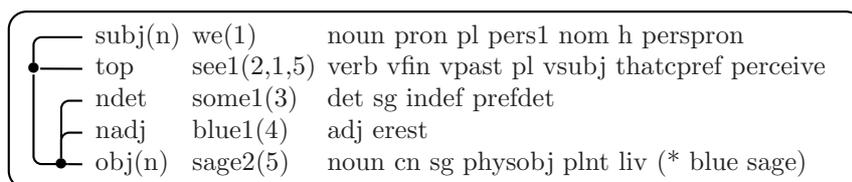
```
=attorney general < n nobj h title
```

By default, a multiword entry, if it applies, is “irrevocable” in the sense of not allowing other analyses of the string it matches. But if a multiword frame has the feature *sepmw* (“separable multiword”), then the parser is allowed to create alternative analyses. For instance the *ESG* lexicon has a *sepmw* entry for “in that” as a *subconj*, and this allows alternative parsing, as needed for “in that house”.

11.2 Structured lexical multiwords

Certain multiwords, which we call *structured multiwords* or *star multiwords*, can exhibit a “have-it-both-ways” scheme, where syntactic structure is shown for the multiword string, but the multiword still behaves basically as a unit, or is “insular”, in parsing: Proper subphrases of the multiword cannot modify, or be modified by, phrases outside the multiword.

Example: If “blue sage” has a lexical entry that makes it a star multiword, then for the sentence “We saw some blue sage”, we get the parse:



The semantic types *physobj plnt liv* come from type markings on the whole multiword “blue sage”, not just the headword “sage”. They come ultimately from the WordNet-based lexical augmentation described in Section 16 below, where there is a star multiword entry for “blue sage” with such types marked. And yet the parse shows syntactic structure of “blue sage”, helping one to know that blue sage is a kind of sage (in a limited sense of “sage”). But note that the parse also shows the string for the multiword, via the feature (* blue sage) of the headword.

The multiwords that can get structured in this way are restricted as follows: They must have exactly one lexical sense, which is a common noun, and the headword is the last word of the multiword. Call these “nouny multiwords”.

Note that “attorney general” wouldn’t satisfy the last condition. Nouny multiwords can be designated as structured by marking a star (*) as a feature in their sense frame.

In the WordNet-based lexical augmentation described in Section 16, approximately 28,000 multiwords in the WordNet-based augmentation get starred. The criteria for marking a multiword W in this way are: (1) W is a common noun and has an **ESG** nouny analysis (**ESG** parsing is used to check this). (2) There are no punctuation marks in W . (3) The components of W are all common nouns, adjectives, or verb participles that are found in the ESG base lexicon.

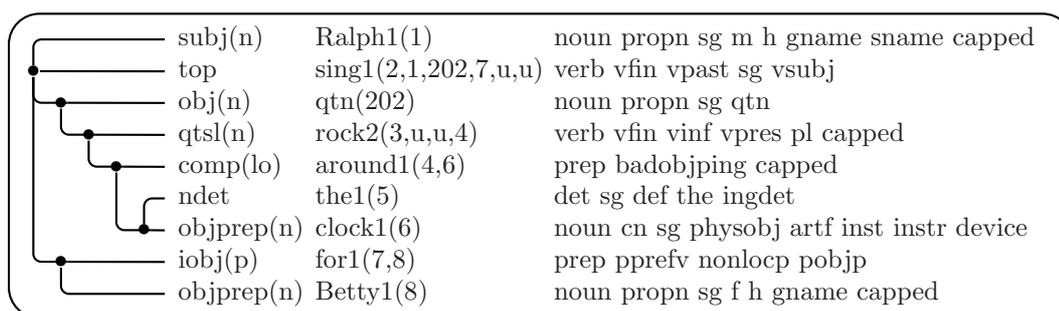
The structuring of star multiwords happens at parse time, but it is done in an efficient way – it doesn’t add significantly to parse time. It is done at a stage before real parsing starts – at the stage where multiwords are being glommed and dealt with. But it does use regular parsing, just for those multiword strings.

There are two flags that control how the multiwords are handled: **parsemw** (on by default), and **parsemwa** (off by default). If both flags are off, then starred multiwords are treated just as normal (non-structured) multiwords. If **parsemw** is on, then all nouny multiwords with a star (*) in their lexical features will get structured. If **parsemwa** is on (the “a” suggests “all”), then *all* nouny multiwords will get structured.

11.3 Quote nodes

Another way **XSG** can deal with chunks, in this case for named entities, is through *quote node*, or *qtnode* analyses in parse trees. Here is an example:

“Ralph sang Rock Around the Clock for Betty.”



Note that there is a proper noun analysis above the subphrase “Rock Around the Clock”, which has sense name **qtn** and feature **qtn**. This is a *qtnode*. This node has one slot **qtsl**, which is filled with a VP analysis of “Rock Around the Clock”. So we “have it both ways”, where the NE is clearly a proper noun chunk, but underneath there is an appropriate syntactic structure.

Qtnode analyses are based not on lexical entries, but are built at syntactic analysis time mainly through attention to capitalization and possible quoting.

The NE may be actually quoted, and that helps, but quotes may be absent, as in the above example, and then capitalization is the main clue. Of course certain “little words” (like prepositions and articles) may be in lower case, also as in this example. But the needs of slot filling also play a role. Since prepositions may be in lower case, what is to prevent the whole string “Rock Around the Clock for Betty” being taken as the qtnode? What happens in this example is that “sing” allows a “for”-PP *iobj*, and the filling of this gets rewarded.

Qtnode analysis is turned on or off by the flag `qtnodes`, which is on by default.

11.4 Named entity lexicons

Another method for dealing with named entities in *XSG* is lexically based. *XSG* can accept a NE lexicon, which should be named `ch.1x`. The `ch.1x` file should follow the syntax of *XSG*/lexicons, as described in [13], with the relaxation that the head word of an entry (normally a multiword) can stand just by itself (with no sense frame information), in which case it is taken to be a proper noun with no other features. But features can be given, after the usual `<` sign that separates head words from sense frames.

The look-up in `ch.1x` can handle very large lexicons efficiently. *ESG* can use a large `ch.1x` NE lexicon which I extracted from Wikipedia, with approximately 1.4 million entries. The extraction routine looks at Wikipedia links instead of titles, because links usually respect case, and so are a better clue to proper noun vs. common noun readings. Various sorts of filtering are done in the extraction, including *ESG* parsing of the candidate entries, to rule out problematic candidates for NEs. The extraction routine is actually bundled with the *ESG* executable, and may be made callable in a subsequent release. But the *ESG* user can construct his/her own version of `ch.1x` and use it with *ESG*.

For using a `ch.1x` lexicon, one must compile it using the command:

```
esg -compilex ch.1x
```

Then to cause *ESG* to read in `ch.1x` upon initialization and use it in parsing, the flag `chlex` should be set on in initialization.

By default, analysis with `ch.1x` just shows the NEs found as multiwords, as in the following example. Note that the multiword has the feature `chn`.

“She sang Sing a Song of Sixpence.”

┌	subj(n)	she(1)	noun pron sg def nom f h perspron
├	top	sing1(2,1,7,u,u,u)	verb vfin vpast sg vsubj
└	obj(n)	Sing a Song of Sixpence(7)	noun propn sg capped chn

However, `ch.1x` analysis can be similar to qtnode analysis in having a “have it both ways” option. If the flag `chstruct` is turned on, then `chn` nodes are still

proper nouns, but show syntactic structure of the multiword underneath. For the preceding example, we then get an underlying VP analysis of “Sing a Song of Sixpence”:

subj(n)	she(1)	noun pron sg def nom f h perspron
top	sing1(2,1,202,u,u,u)	verb vfin vpast sg vsubj
obj(n)	Sing a Song of Sixpence(202)	noun proprn sg chn
chsl(n)	sing1(3,202,5,u,u,u)	verb vfin vinf vpres pl capped
ndet	a(4)	det sg indef
obj(n)	song1(5,u,6,u)	noun cn sg abst artcomp musicomp comm
nobj(n)	of1(6,7)	prep pprefn nonloep
objprep(n)	sixpence1(7)	noun cn sg pl sgpl abst comm

11.5 Creating chunks at tokenization time via tags

Normally, text tokenization by *XSG* follows the usual rules for tokenization, paying attention to runs of letters, delineation by punctuation symbols, effects of abbreviations, and the like. But some special *XSG* tag pairs – let us call them “chunking tags” – can be used to delineate larger tokens, or “chunk” tokens, so that any characters, even blanks and newlines, can appear in the chunk (aside from the end-tag of the tag pair).

Chunking tags can be used to delineate named entities, making multiword (or single-word) named entities into chunks, and even specifying semantic types for the chunks, which will figure in the parsing. Such tagging of named entities might be done by a preprocessor for *XSG*, which passes the tagged text to *XSG* for better parsing. Actually, the *XSG* chunking tags can be used to provide much more general lexical specifications for chunks than are normally considered for “named entities”, as we will see below.

In the next subsection we will describe an additional method for providing “chunking control” for *XSG* at tokenization time. Instead of using tags, this method involves function calls to a hook (call-back) function that is monitoring the text for the existence of chunks at particular positions in the input document.

There are two classes of chunking tags known to *XSG*: (a) *literal* or *delimiter* tags, and (b) *lexical* tags.

When a chunk *xxx* is surrounded by a *literal* tag `<ttt>` (and its end tag) as in

```
<ttt>xxx</ttt>
```

the chunk *xxx* will be treated like any other word that is looked up in the *XSG* lexicon (with morphology) for parsing. If the chunk is not found in the lexicon, then it will be given the default lexical analysis for all unfound words – a singular proper noun. Currently, *XSG* knows only one literal chunking tag, named *lit*.

The basic idea of a *lexical* tag `<ttt>` is that some of the lexical and semantic feature control lies with the tag itself. The chunk `xxx` surrounded by the tag and its end-tag may or may not be actually looked up in the lexicon, but even if it is looked up, the tag may filter out lexical analyses and may specify additional semantic features of the chunk. What is done depends on the attributes within the tag `<ttt>`, to be described below.

Three particular lexical tags – `<ch>` (for “chunk”), `<ne>` (for “named entity”), and `<pn>` (for “proper noun”) – can contain attributes that specify lexical features of the chunk `xxx`. In spite of the different names for these three tags, they are all treated the same by *XSG*. You might as well use `<ch>` in all cases. I have kept `<ne>` and `<pn>` for compatibility reasons. In the discussion below, I will use only `<ch>`, but you could use either of `<ne>` and `<pn>` as well.

Here is an example of use of a lexical tag:

```
<ch lex="proprn Company">International Business Machines</ch>
```

This will make “International Business Machines” into a chunk which is a proper noun (`proprn`) and has the semantic type `Company`.

There are two types of attributes a lexical tag can have – `lex` and `sem`, which should be used separately (not both together in the same lexical tag).

Let us look first at `lex`. With the `lex` attribute, one completely specifies the lexical analysis of the surrounded chunk `xxx`, and `xxx` will not be looked up in the lexicon. In general, the form of a `lex` attribute-value pair in a `<ch>` tag is:

```
lex="LexicalAnalysis"
```

Note: No spaces are allowed on either side of the = sign. The same is true for other attribute value pairs below. The `LexicalAnalysis` is in general what can appear as a lexical analysis in an *XSG* lexicon (see [13]). (There is an exception in the notation, described below.) Such lexical analyses can specify not only proper noun frames, as in the above example, but frames of any part of speech, and even more than one frame.

Here is a more complex example:

```
<ch lex="n (p for) ~~ v obj">fixem up</ch>
```

This would give the multiword `fixem up` both a noun analysis `n (p for)` (a common noun with a `for`-prepositional complement slot) and a verb analysis `v obj` with a direct object slot. The symbol `~~` is used to separate lexical analyses in the `lex` attribute instead of the usual `<` in *XSG* lexicons. These lexical analyses can express anything that normal *XSG* lexical entries can – any parts of speech, slot frames, semantic types, subject area tests, etc.

One constraint, when the `lex` attribute is used, is that no morphology is done on the chunk. If the chunk appears in an inflected form, then the `lex` attribute must itself specify the morphological features. So for a noun plural `fixem ups`, you might write something like:

```
<ch lex="n (p for) pl">fixem ups</ch>
```

thus showing the noun plural feature `pl`.

Actually, for the LMT machine translation system, the `lex` attribute can show even more than just described. It can show complete LMT bilingual lexical entries for chunks, and LMT will use those in parsing (with *XSG*) and translation. In the `lex` attribute notation, the symbol `$$` is used instead of the symbol `>` that is normally used in LMT lexicons to introduce transfers.

Now let us look at lexical chunking tags where the `sem` attribute is used. Here is an example:

```
<ch sem="n PERSON MALE">fathers</ch>
```

When the `sem` attribute is used, the surrounded chunk *is* looked up in the lexicon, with morphology. Initially, in this example, the lexical analysis shows `fathers` both as a plural noun and as a third person singular present tense verb. However, with `sem`, only one lexical analysis will be used, which in this example will be the noun analysis, because of the initial POS symbol `n` in the value of `sem`. In addition, the semantic features `PERSON` and `MALE`, included in the value of `sem`, are added to the lexical analysis, and they will show in the parse tree for this node.

In general, the value of the `sem` attribute should be a string of the form

```
[POS] SemFea1 SemFea2 ...
```

The optional part of speech POS should be chosen from these part of speech symbols:

```
n, propn, pron, num, v, adj, adv,  
det, prep, subconj, conj, qual
```

These are parts of speech that can appear in *XSG* lexical entries. The first four are used for subcategories of nouns – common noun, proper noun, pronoun and number, respectively. The meanings of the others should be clear. If the POS is omitted, the default `propn` is used. The subsequent items `SemFeai` in the value of `sem` are semantic features, and will be added to the semantic features of the lexical analysis. These features need not be “known” to *XSG*. They will appear in the parse tree on the node corresponding to the chunk.

So, given a `ch` tag with a `sem` attribute, *XSG* will do morphology and lexical look-up on the surrounded chunk. But *XSG* will select only one lexical analysis. If there is an analysis whose POS matches that of the `sem` tag, then *XSG* uses the first such. Otherwise, if the POS of the `sem` is `propn` and there is an `n` (common noun) lexical analysis, *XSG* will use the first such. Otherwise, *XSG* will take the analysis to be `propn`. In all cases, any semantic features listed in the value of `sem` will be added to the lexical analysis.

If the `<ch>` tag has no `lex` or `sem` attribute, as in

```
<ch>International Business Machines</ch>
```

then the chunk will be given the default lexical analysis of a singular proper noun.

So there is a great deal of generality in the lexical specification allowed with the `<ch>` tag; it includes more than is usually encompassed under the notion of “named entity”. I wish to thank Nelson Correa for a useful discussion about using *XSG* chunking tags to specify properties particularly of named entities, which led me in part to the general formulation described here with lexical chunking tags.

XSG knows several other lexical chunking tags. When HTML tag reading is turned on, the additional tags of this sort (besides `<ch>`, `<ne>` and `<pn>`) that are recognized are:

```
<code>, <samp>, <var>, <l>
```

XSG does not recognize `lex` or `sem` attributes in these, and the surrounded chunks just get the default lexical analysis of singular proper noun.

11.6 Creating chunks at tokenization time via call-back functions

Now let us look at the second way of specifying chunks and their lexical analyses for *XSG* at tokenization time (not using tags). This involves a “chunk hook function” (defined by the user of *XSG*) which knows about document positions and tells *XSG* where the chunks start and end (in terms of document positions), and what their lexical analyses are. This methodology is used when *XSG* is embedded in UIMA (see the Introduction). I wish to thank Marshall Schor for discussions that led to this formulation.

To explain the form and use of chunk hook functions, I need to say a bit about the way *XSG* handles document positions.

When *XSG* is processing a file, the system keeps track of line numbers and column numbers in the file. Column numbers are like byte offsets on a line, so a column number of 0 is used for the position right before the first character. (You can tell *XSG* to process a file either by calling the API function `xSGdoFile` described below in Section 14, or by using the command `do` in interactive mode, as described above in Section 5.) So in the case of file input, document positions are specified by ordered pairs (*LineNo*, *ColumnNo*), and the hook function needs to deal with these.

On the other hand, *XSG* can process a document given to it directly as a string, with the *XSG* output also produced as a string. (You can tell *XSG* to do this via the API function `xSGdoSt`, described below in Section 14.) In this case, all that counts for giving a document position is the byte offset in the document string. Newlines are treated as normal characters that count in

the byte offset. But *XSG* still deals formally with line numbers and column numbers in this scenario, and the chunk hook function must also. The “line number” is always just kept at 0, and the “column number” is the byte offset in the whole document string.

Given this explanation of line numbers and column numbers for specifying document positions, we can now describe the form of a chunk hook function for specifying chunks. Such a function (let’s call it *chhook* here) should have the following type:

```
sint chhook(ptr id, sint slineno, sint scolno,
            sint *elineno, sint *ecolno, st *annotation)
```

Here the types are defined by:

```
typedef void * ptr;
typedef long int sint;
typedef unsigned char * st;
```

(The type *st* has a different definition in case *XSG* is compiled in Unicode mode – see Section 12.) The arguments of *chhook* are as follows:

1. *id* is the SG handle, as described in Section 14. You obtain this by calling the SG initialization function *xSGstart*, described in that section.
2. *slineno* and *scolno* give the starting line and column number pair of the chunk. These are passed by the *XSG* tokenizer to *chhook*.
3. **elineno* and **ecolno* give the ending line and column number pair of the chunk. These are passed back to the tokenizer by *chhook*.
4. **annotation* is a string that is an attribute-value list. This is returned by *chhook* to *XSG*. This list can contain one of the attribute-value pairs

```
lex="LexicalAnalysis"
sem="SemAnalysis"
```

of just the same form allowed as *lex* or *sem* attribute-value pairs on the chunking tag *<ch>*. Then *XSG* will associate a lexical analysis with the chunk, in just the same way it would if *<ch>* had been used. Actually, *XSG* stores the whole **annotation* list with the chunk, but currently *XSG* does not use any more of it than a *lex* or *sem* attribute-value.

During tokenization, *XSG* calls the *chhook* function at every spot where a possible new token could begin. It passes in the current document position to *chhook* in the pair (*slineno*, *scolno*). If *chhook* *succeeds*, that is, returns a non-zero integer, then *XSG* will create the chunk from all characters in the document stretching from the current position up to the ending position

(`elineno`, `ecolno`) returned by `chhook`, or up to the end of the document, whichever comes first. And it will associate the `*annotation` list with the chunk, as indicated.

In order for `XSG` to call a chunk hook function `chhook` during tokenization, you have to make `chhook` known to `XSG` as such, and you do this by calling the API function

```
void xSGsetChhook(ptr id,
                 sint (*chhook)(ptr, sint, sint,
                               sint *, sint *, st *))
```

where the first argument `id` is the SG handle and the second argument is (a pointer to) the hook function. So if the hook function is named `chhook`, you would call:

```
xSGsetChhook(id, chhook);
```

Figure 6 shows a simple example for `chhook`, used with the file-mode of processing (where the line number coordinates vary), which for each line will cause the string stretching from column position 3 to column position 14 to be a common noun with semantic type `Weird`.

```
sint chhook(ptr id, sint slineno, sint scolno,
           sint *elineno, sint *ecolno, st *lex) {
  if (scolno != 3)
    return 0;
  *elineno = slineno;
  *ecolno = 14;
  *lex = "lex=\n Weird\>";
  return 1;
}
```

Figure 6: Example of a chunk hook function

So what would this chunk hook function do for the following text?

```
In tagged text, tag pairs within segments like those that
indicate font changes, italics and boldface are often useful
to XSG in delineating subphrases of segments.
```

First `chhook` is called at position (1, 0), because a token is starting there, but `chhook` fails. It is next called at position (1, 3), and here it succeeds, creating the chunk `"tagged text"` (up through position (1, 14)). It is called several times more on the first line, but without success. On the second line, the only

chance for a successful call would be at position (2, 3), but the tokenizer is in the midst of doing the token "indicate" at that point, and no call is made. On the third line, there is a successful call at position (3, 3), which creates the token "XSG in deli" (through position (3, 14)). The next token on that line (recognized by normal tokenization, not by `chhook`) begins immediately, and is "neating".

We have discussed two ways of specifying chunks of text for *XSG* and giving them lexical analyses. There is a third way that will work for most kinds of chunks – simply to put them in a user lexicon for *XSG*. User lexicons are described below in Section 15, and they use the general SG lexical format specified in [13]. This will work in many cases, but there are advantages to using chunking tags or a chunk hook function:

- Named entities may be determined on-the-fly by a named entity recognizer that uses general rules instead of depending on a known list.
- Chunks (as specified in this section) can consist of *any* characters, whereas multiwords in *XSG* lexicons are more restricted.
- If an *XSG* lexicon has too many multiwords (e.g. in the hundreds of thousands), the processing may not be efficient enough. (This problem does not exist for single-word entries.)

12 The data structure for SG parse trees

The *XSG* API described in this report uses a certain data structure for parse trees that we call `sgph` (for “Slot Grammar phrase”). We describe this data structure in the current section. In the next section, we describe the representation of punctuation and tags in *XSG* analysis structures.

The `sgph` data structure is considerably simpler than the data structure for phrases that SG uses internally during parsing. The definition of `sgph` involves only a total of six datatypes (including itself), but the definition of the internal phrase structure involves about eighteen datatypes. The external API datatype `sgph` contains essentially all the same applicable information as the internal one, but the internal one has finer distinctions for the sake of efficiency and control of the parsing process itself. For instance, as mentioned in Section 7, grammatical features are encoded via bit string positions in the internal representation.

In the `sgph` structure, features, slots, and slot options are represented as strings. The string datatype is named `st`. For the single-byte case, `st` is defined by:

```
typedef unsigned char * st;
```

For the double-byte (Unicode) case, `st` is defined by:

```
typedef wchar_t * st;
```

The list of features for an `sgph` structure is represented by a string list, `stlist`, defined by:

```
typedef struct Stlist {st hd; struct Stlist * tl;}
    * stlist;
```

We also find it convenient to have a separate name, `sint`, for long signed integers, defined by:

```
typedef signed long int sint;
```

And we define `ptr` (for “pointer”) by:

```
typedef void * ptr;
```

Now we can define the phrase structure `sgph` itself. It is given in Figure 7. To be complete, this requires only the definition of `struct Sgphlist` (and the type of a pointer to it), which is given in Figure 8.

```
typedef struct Sgph {
    sint wordno; sint lb; sint rb;
    stlist f;
    st word; st lword; st cite;
    st sense; st ssense;
    intlist osenses; dbllist oweights;
    stlist compslots; struct Sgphlist * frame;
    struct Sgphlist * lmods; struct Sgphlist * rmods;
    st sl; st opt;
    struct Sgph * mother;
    double eval;
    ptr ref;
} * sgph;
```

Figure 7: Definition of the `sgph` data structure

```
typedef struct Sgphlist {
    sgph hd;
    struct Sgphlist * tl;
} * sgphlist;
```

Figure 8: Definition of the `sgphlist` data structure

Notice that these two data structures are mutually recursive. If we express them totally in terms of the pointer types, we see that an `sgph` is a pointer to a structure with the fields:

```

sint wordno; sint lb; sint rb;
stlist f;
st word; st lword; st cite; st sense; st ssense;
intlist osenses; dbllist oweights;
stlist compslots;
sgphlist frame;
sgphlist lmods; sgphlist rmods;
st sl; st opt;
sgph mother;
double eval;

```

And a phrase list, or `sgphlist`, is a pointer to a structure with the two fields:

```

sgph hd;
sgphlist tl

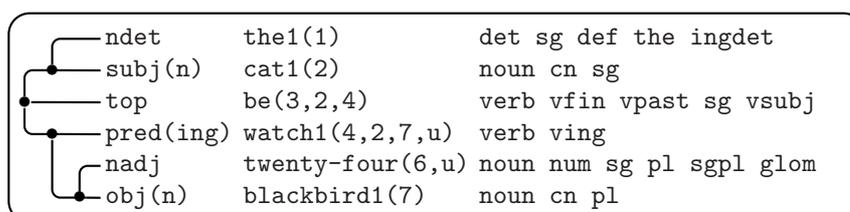
```

It is best to view `sgph` and `sgphlist` that way as we look further at what these structures represent.

The idea of the data structure `sgphlist` should be clear; it is just a list of `sgph` objects. So it remains to describe `sgph`.

Let us go through the different fields of an `sgph` and describe what they mean. We will look at an example parse tree to illustrate each point, the parse for the sentence:

The cat was watching twenty-four blackbirds.



So here are the individual fields of an `sgph` and what they mean:

- `wordno`. The word number in the sentence of the head word of the phrase. In our example, for the subphrase for *the cat*, it is 2. This number shows in the parse tree display as the first argument of the word sense predication.
- `lb`. The left boundary of the phrase in the sentence. It is one less than the sentence word number of the leftmost (single) word in the phrase. In our example, for *twenty-four blackbirds*, it is 4. In this example, the tokenizer sees *twenty* and *four* as separate tokens, and the word number of *twenty* is 5.
- `rb`. The right boundary of the phrase in the sentence. It is the sentence word number of the rightmost (single) word in the phrase. For *the cat*, it is 2.

- f. The feature list of the phrase, as a list of strings. For the whole sentence in the example, then, it is the list consisting of:

verb vfin vpast sg vsubj

word. The head word, exactly as written in the input, including capitalization. For the **ndet** filler, it is **The**.

lword. The lower case form of the head word. For the **ndet** filler, it is **the**.

cite. The citation form (after inflections are “undone”) of the head word. For the **obj** filler, it is **blackbird**.

sense. The sense of the head word. For the **obj** filler, it is **blackbird1**. For *twenty-four*, it is **twenty-four**. In parse tree displays, the default is to show the **sense** string as the predicate portion of each word sense predication. One can see this in the display of our example. However, flags can be set to show either the **cite** field of the **ssense** field instead; see Section 9, page 36.

ssense. A variant of the sense of the head word. In most cases, it will be exactly the same as **sense**, but for example for literal numbers, a numeric representation of the number is shown. For *twenty-four*, the **ssense** looks roughly like this:

(lnum 24.00)

For ordinal numbers like *twenty-fourth*, the **lnum** would be replaced by **olnum**.

osenses, oweights. These have to do with word sense disambiguation, and we omit the discussion here.

compslots. The list of names of the complement slots of the head word sense. The list is given in the order of the complement slots specified in the lexical entry for the word sense, or derived from such a lexical frame.³

frame. This is the list of filler phrases of the complement slots of the head word sense. They are in 1-1 correspondence with the slots (slot names) in the **compslots** field. I.e., the *n*th member of **frame**, if it is non-null, is the filler of the slot named by the *n*th member of **compslots**. The **frame** field is reflected in the arguments of the word sense predication in the parse tree – the arguments after the first one (which is the **wordno**). This is explained in part on page 5 above in Section 2. In our example, the word sense predication for *watching* is:

³Some complement slots may not be required in lexical entries (and so are implicit); for example the **subj** slot of a verb need not be listed, unless it has special conditions on its fillers. A word sense may be obtained by derivational morphology; in this case the morphology will create a derived complement slot frame from that of the base word.

```
watch1(4,2,7,u)
```

In this form of the display, each member of the **frame** list is exhibited only by its **wordno**. The 2 is for the logical subject of *watching*, which is *the cat*. The 7 is for the logical direct object, which is *twenty-four blackbirds*. If a slot of the complement slot frame is not filled, then the corresponding member of **frame** is **nil** (the null pointer), and is indicated in the parse display by **u** (“unfilled” or “unknown”) in the word sense predication. If the flag **predargslots** is turned on (see Section 9), then the slot from **compslots** corresponding to each **frame** member is also shown in the parse display, like so:

```
watch1(4,subj:2,obj:7,comp:u)
```

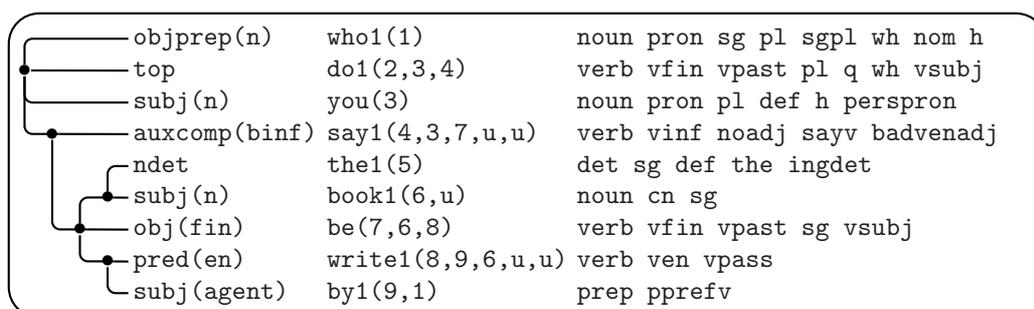
The slots and fillers in **compslots** and **frame** are the *logical* slot-fillers for the word sense. For example a **subj** slot is the *logical* subject.

- lmods**. The list of left modifiers of the head word (each member of which is an **sgph** object). They appear in left-to-right sentence order in the list. Both they and the right modifiers are shown in the obvious way in the tree display.
- rmods**. The list of right modifiers of the head word. They also appear in left-to-right sentence order in the list. In the internal phrase data structure, the right modifiers are shown in the reverse order of this, for efficiency reasons.
- s1**. The slot filled by the phrase. This is shown in the parse tree display just to the left of the word sense predication.
- opt**. The option chosen for the slot filled by the phrase. For (most) complement slots, this is shown in parentheses after the **s1** field. It will be shown also for adjunct slots if the flag **showaopts** is turned on (see Section 9). For the top-level phrase, the slot is **top**.
- mother**. The mother phrase (again an **sgph**) of the phrase. It is the null pointer for the top-level phrase, but is non-null for all subphrases.
- eval**. This floating point number (a **double**) gives a parse score for this phrase. We will not try to describe parse scoring in this report, but one thing that should be noted is that a higher score is a worse score.

Note that the **sgph** parse tree is “well-pointered”. There is the upwards pointer provided by the **mother** field, and the downwards pointers in the **lmods** and **rmods** fields. Also there are the pointers in the **frame** field, which can point to a remote filler, as in examples like

Who did you say the book was written by?

where the parse tree is



Note here that the second argument (with `wordno` equal to 1) of the `by1` node points back to the `who1` node, showing that *who* is the `objprep` of *by*. In turn the agent `by1` node (with `wordno` equal to 9) is shown as the second argument of the `write1` node, which is for the logical subject of `write1`. This means essentially that *who* is the logical subject of *write*.

13 Punctuation and tags in SG analyses

The `sgph` data structure described in the preceding section by default does not show punctuation symbols or tags that appear in the input sentence. Many punctuation symbols, such as commas that set off adverbial modifiers of clauses, or font-change tags, are not crucial to the grammatical structure of the sentence. The `XSG` parser *does* notice them and use them in constraining parses; but they are often optional to the grammaticality of the sentence, and their presence in the parse output is not essential to some applications. In fact, the omission of them in the parse structure simplifies most applications of `XSG`, because this reduces the number of cases one has to look at in postprocessing of the parse structures. However, they are important for some applications, and they are indeed accessible in the `XSG` API. We describe this in the current section. More generally, we describe how to see the results of `XSG` tokenization.

The kinds of tags that can be omitted in the parse trees include font-change tags, HTML anchor tags, the chunking tags described in Section 11, tags that denote punctuation symbols, and others. Such tags usually occur in pairs (begin and end tag) within a segment, and often (but not always) bracket grammatical phrases. The tag pairs can sometimes cross segments, however. We call these *bracketing* tags in general, also sometimes *highlighting* tags.

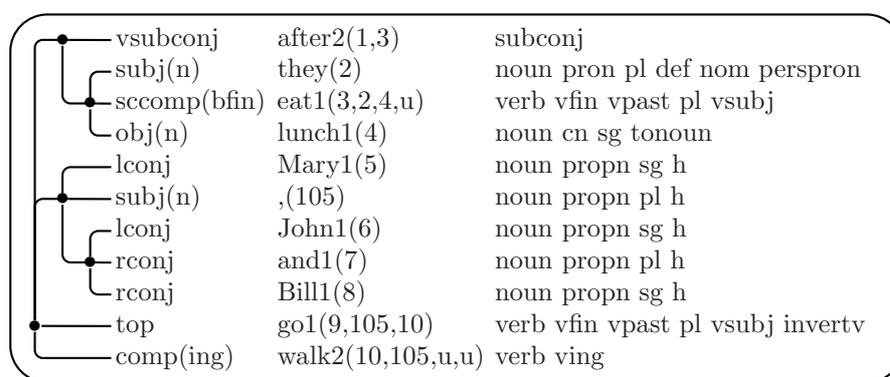
Of course some tags are used to denote grammatical elements that act as words or whole phrases in the segment, especially proper nouns. These *are* shown in the parse tree.

In the remainder of this section, we will usually just use the term “punctuation” to include both (normal) punctuation and bracketing tags.

There is an exception to the omission of punctuation in the *XSG* parse trees. When a punctuation symbol plays the role of a coordinator, then it is shown in the parse tree. Consider for example the sentence:

After they ate lunch, Mary, John, and Bill went walking.

The parse tree is:



Note that the parse tree shows only one comma – the one after *Mary* that acts as a coordinator. (This is the case for the *sgph* data structure also.) The first comma (after *lunch*) is merely a phrase separator, and could even be omitted without loss of grammaticality in the sentence. The third comma is of course connected with coordination, but *XSG* considers the real coordinator there to be *and*, with the comma used in an ancillary way. (Note that this comma is optional also.) The period at the end of the sentence is also not reflected in the parse tree.

When a punctuation symbol, like the comma after *Mary* in our example, is used as a coordinator, we say that it is *promoted* to a sentence tree node. In such a case, the promoted symbol is assigned a node number which is 100 plus the node number of the preceding word. In our example, *Mary* is the 5th word, so the comma after it gets node number 105 in the parse tree. In counting *words* of the sentence, one does not count the punctuation tokens. The constant 100 is actually a #defined constant in *XSG*, called `sentlenmax`, which sets a maximum on the number of words in a segment that *XSG* will try to parse. *XSG* can be compiled with a higher value for `sentlenmax`.

So how does one see the punctuation symbols that are *not* promoted? The intuitive answer is that they are stored “in the interstices” between the words of the segment. The *interstices* are indexed by integers that run from 0 to the word number for the last word of the segment. The interstice indexed by 0 holds the punctuation, as a list of tokens, that occurs before the first word. For $i \geq 1$, the interstice indexed by i holds the punctuation occurring right after word i . Punctuation symbols that are promoted are also shown in these token lists in the interstices, in their proper place.

Using `Actions`, you can open the data base.

type	id	word	lword	spword	tagname	sp
wordtok	gtok	"Using"	"using"	"Using"	"using"	{3, 0, 3, 5}
lhili	gtok	""	""	" "	"b"	{3, 6, 3, 9}
wordtok	gtok	"Actions"	"actions"	"Actions"	"actions"	{3, 9, 3, 16}
rhili	gtok	""	""	""	"/b"	{3, 16, 3, 20}
septok	comma	","	","	","	","	{3, 20, 3, 21}
wordtok	gtok	"you"	"you"	" you"	"you"	{3, 22, 3, 25}
wordtok	gtok	"can"	"can"	" can"	"can"	{3, 26, 3, 29}
wordtok	gtok	"open"	"open"	" open"	"open"	{3, 30, 3, 34}
wordtok	gtok	"the"	"the"	" the"	"the"	{3, 35, 3, 38}
wordtok	gtok	"data"	"data"	" data"	"data"	{3, 39, 3, 43}
wordtok	gtok	"base"	"base"	" base"	"base"	{3, 44, 3, 48}
termtok	dot	."	."	". "	."	{3, 48, 3, 49}

Figure 9: Display of token analysis for a segment

The next step in explaining these notions is to make precise what a *token* is. The *XSG* tokenizer produces, for each segment, a list of objects of type `token`. These include both the *words* of the segment and the punctuation between the words. One can see a display of the tokens for a segment by turning on the flag `toktrace`. An example of `toktrace` output for a sentence is shown in Figure 9. (We have lined up the output in a tabular form that is a bit easier to read.)

In Figure 9, seven fields of each token are displayed; they are labeled at the beginning of the output. There is an eighth field, called `annotation`, which is not shown in `toktrace` output, and we explain that below. Let us describe the seven fields.

The `type` field of a token (of C datatype `sint`) shows the general category of the token. The most common type is `wordtok`, for normal words. The type `lhili` (for “left highlight”) is for begin bracketing tags, and `rhili` is for the corresponding end tags. The type `septok` is for punctuation (like commas and dashes) that can serve as a separator between a modifier and its modifiee, but such a token may also be promoted to a coordinator. The type `termtok` is for segment terminators. We will not try to list all the token types in this report, but the API function

```
st xSGtoktype(token tok)
```

provides the string form of the `type` field of token `tok`.

The `id` field (of C datatype `sint`) is used to name very special tokens that figure heavily in the *XSG* grammar. In Figure 9 one sees three examples, `comma` and `dot`, with obvious meanings, and `gtok` (standing for “general token”). There are quite a few possible `ids` in the system, mainly corresponding to specific punctuation symbols, and we will not list them here. The API function

```

typedef struct { sint slineno; sint scolno;
                sint elineno; sint ecolno; }
    * span;

typedef struct { sint type; sint id;
                st word; st lword; st spword; st tagname;
                st annotation; span sp; }
    * token;

typedef struct Tokenlist {token hd; struct Tokenlist * tl;}
    * tokenlist;

```

Figure 10: Definition of the `token` data structure

```

st xSGtokid(ptr handle, token tok)

```

provides the string form of the `id` field of token `tok`. The first argument is the SG handle, as described in Section 14.

All the rest of the fields are of datatype `st` (string), except for `sp`.

The `word` field is the underlying string for the token, just as it appears in the text, including capitalization.

The `lword` field is like the `word` field, but is made lower case.

The `spword` field (for “space plus word”) consists the `word` string, prefixed by any whitespace characters that occur between the preceding token and the current token. One can reconstruct the whole segment string by concatenating all the `spword` fields of the tokenlist created by the tokenizer.

For tags, the `tagname` field is the name of the tag, i.e. the part of the tag string that starts after the initial `<` and goes up to a blank or otherwise up to the closing `>`. It is also normalized to lower case. For non-tags, the `tagname` is the same as the `lword`.

Finally, the `sp` field of a token gives its *span*. This represents the start and end positions of the token in the input document (not counting preceding whitespace). The `sp` field is of type `span`, defined in Figure 10. It is a pointer to a quadruple of integers giving, respectively, the start line and column numbers and end line and column numbers of the token. See the discussion on page 44 in Section 11 for more detail on how these positions are calculated.

The *XSG* data structure for a `token` is defined in Figure 10. In this figure, we define not only `token` but also the auxiliary data types `span` and `tokenlist`.

From the preceding discussion, it should be clear what the meaning is of all these datatype ingredients, except for the `annotation` field of a `token`. This relates to chunking, as described in Section 11. If chunking is done through a chunk hook function (see page 44), then the `annotation` field of the chunk

token is set to the annotation returned in the last argument of the chunk hook function.

Now let us return to the way the *XSG* API allows one to see punctuation and word tokens.

First, the API function

```
token * xSGwds(ptr id)
```

where *id* is the SG handle (as described in Section 14), returns an array of the regular word tokens for the most recently parsed segment. For convenience, let us call this array *Words* in the following. The API function

```
sint xSGwdsLen(ptr id)
```

returns the index of the last word of the segment. Let us denote this by *n* in the remainder of this section. Thus the the members of *Words* are indexed by 1 through *n* inclusive.

Next, the API function

```
tokenlist * xSGpuncs(ptr id)
```

returns an array – call it *Puncs* – such that for each $i, 0 \leq i \leq n$, *Puncs*[*i*] is the list of punctuation tokens in interstice *i*.

For the example in Figure 9, the arrays *Words* and *Puncs* would look like this:

i	0	1	2	3	4	5	6	7	8
<i>Words</i> [<i>i</i>]		Using	Actions	you	can	open	the	data	base
<i>Puncs</i> [<i>i</i>]	()	()	(,)	()	()	()	()	()	(.)

In the *Words* row we are displaying the word tokens by their *word* fields. In the *Puncs* row we are displaying the punctuation tokenlists in the form $(w_1 w_2 \dots)$, where w_j is the *word* field of the *j*th list member.

It should be noted that the tokens in the *Words* array do not show multiword agglomeration. For example, the *ESG* lexicon has a multiword entry for *data base*. In the *ESG* parse of our example sentence, *data base* will show as a single node. Its *wordno* field will be 8, and no node in the tree has *wordno* 7. In general, the parse *wordno* of a multiword will be the *Words* index of the token that is the head word of the multiword.

14 Compiling your own XSG application

If you have the object files or a DLL for the *XSG* modules and you want to compile your own application that uses *XSG*, then the following information is relevant.

You will get a header file `xsg.h` with the SG datatypes given in Section 12, and the prototypes for the API functions described in the current section. And you will get a DLL `Xsg.dll` and `Xsg.lib`.

In the following, let us use the following abbreviations (these definitions are in `xsg.h`).

```
#define nil NULL
typedef FILE * file;
```

Two API functions that you will need to call in your own program are:

```
sint xSGstart(int argc, unsigned char **argv, ptr *id)
void xSGend(ptr id)
```

The function `xSGstart` should be called to initialize `XSG`. Its first two arguments are like the two arguments of `main` in a C program used to hold command-line arguments of the corresponding executable. Thus `argv` should be an array of strings, and `argc` is the length of that array. These arguments can be used to control initialization of `XSG` and to cause `XSG` to perform certain auxiliary operations. The initialization command-line arguments for the executable `XSG.exe` described at the end of Section 3 can all be given to `xSGstart` in its `argc` and `argv` arguments. We also discuss other aspects of this in Section 15. It is permissible to call `xSGstart` with `argc` equal to 0 and `argv` equal to `nil` if no auxiliary operations are desired. The third argument `*id` will be set by `xSGstart` to a pointer called the SG *handle*, which is used in other API functions. If `*id` is set to `nil`, this means that initialization was not successful – perhaps because of insufficient memory or because `XSG` could not access its lexical files. In this case, `XSG` will write an error message to `stderr`, and `xSGstart` will return 0. The function `xSGstart` may also return 0 if its first two arguments initiate certain auxiliary operations, as discussed in Section 15. Otherwise, 1 is returned.

We mention here one kind of parameter pair that `xSGstart` can take in its first two arguments. A pair of parameters of the form

```
-lexpath Path
```

will allow the `XSG` lexical files to be in another directory than the current directory. The `Path` string should show either the complete directory path (with an initial slash) for the lexical files, or else the subpath (without an initial slash) from the current directory. (The first option would be the normal one.) The `Path` string should not contain a final slash.

You should call `xSGend(id)`, where `id` is the SG handle set by `xSGstart`, when you are finished using `XSG`. This frees storage.

In the following, we will use `id` to refer to the SG handle set by `xSGstart`.

There are two basic ways of using `XSG`:

1. Your program is in charge of file handling and text segmentation, and you pass segment strings to *XSG*, getting parses (*sgph* data structures) back.
2. *XSG* is in charge of processing input texts and files, and calls a “hook” after each parse that invokes a function of yours that does what you want with the parse.

Let us look at these two methods.

- (1) For each segment string *s*, you should call the *XSG* API function:

```
sgph xSGpars(ptr id, st s)
```

This returns the best *XSG* parse as an *sgph* tree.⁴ Note: *xSGpars*(*id*, *s*) should not be called when the number of tokens in *s* is greater than 100.

- You should call the API function

```
void xSGclear(ptr id)
```

sometime after each parse but before calling *xSGpars* the next time. This clears storage space for the next parse. The storage area used in the returned parse from *xSGpars* will get overwritten after you call *xSGclear*, so you should make sure to do what you need to do with the returned parse (like writing results to a file, or storing information in some other storage area of your own).

- Under method (1) there is an alternative API function:

```
sgphlist xSGparsl(ptr id, st s)
```

This returns the list of *all* the parses for segment *s*, ordered according to parse score (best first).

- (2) In the second method of interfacing *XSG* to your own program, *XSG* is in charge of file handling, segmentation, tag handling, etc. You need to define a “hook” function, with a prototype like this:

```
void my_hook(ptr id, sgph ph)
```

After each parse, *XSG* will call *my_hook* with *id* as the *SG* handle and *ph* as the best parse for the sentence. But in order to communicate to *XSG* that this is your hook function, you should make the call:

```
xSGhookfun(id, my_hook);
```

in your program at some point after your call to *xSGstart*.

- There are three ways that your program can invoke the operation of *XSG*.

- (a) A call to

⁴Best according to parse score in the field *eval* of *sgph* phrase structures.

```
void xSGinteractive(ptr id)
```

will put *XSG* into interactive mode, with the “Input sentence” prompt, described in Section 3. As indicated in that section, you can leave this loop by typing “stop.”.

(b) You can call:

```
sint xSGdoSt(ptr id, st text, st *analysis, sint *alen)
```

This takes input text as a string *text*, which may consist of any number of segments. The results (parse displays or other output) are written into the string **analysis*. The length of the output string in bytes is put into **alen*.

(c) A call to

```
sint xSGdoFile(ptr id, st InFile, st Outfile)
```

will process the input file *InFile*, writing results to output file *OutFile*. Actually, the input argument string *InFile* for *xSGdoFile* can be a *pattern* of files (using *** as wildcard), for example *d:\texts*.html*. In this case, all the files matching the pattern will be processed, and all the output goes to the single file *OutFile*.

The API function

```
file xSGout(ptr id)
```

returns a pointer to the (*current*) *SG output file*. If *xSGdoFile* is operating, then the SG output file is the output file specified for that function. If *xSGinteractive* is operating, then the SG output file is *stdout* by default, but will be *sg.out* if the flag *xout* is on and *XSG* is writing results. You can use *xSGout* to write your own data to the SG output file. You can also open the SG output file for a file name *fln* by calling:

```
sint xSGopenOut(ptr id, st fln)
```

You can close the current SG output file by calling:

```
sint xSGcloseOut(ptr id)
```

This will not only close the file, but will also set the current SG output to *stderr*.

Three output functions provided in the API are these:

```
void xSGprst(ptr id, st x)
void xSGpnl(ptr id, st x)
void xSGpri(ptr id, sint x)
```

The first writes string `x` and the second is similar but adds a newline. The third write the integer `x`. “Writing” is either to the current output file in the case of the function `xSGdoFile` or the output string `*analysis` in the case of `xSGdoSt`.

See Section 13 for a description of the following API functions, which allow access to the results of *XSG* tokenization and the treatment of punctuation and tags.

```
token * xSGwds(ptr id)
sint xSGwdsLen(ptr id)
tokenlist * xSGpuncs(ptr id)
st xSGtokid(ptr id, token tok)
st xSGtoktype(token tok)
```

A useful API function for setting flags is:

```
sint xSGsetflag(ptr id, st flg, sint val)
```

This sets a flag with name `flg` to the integer value `val` and returns 1, if `flg` is indeed a valid flag name. Otherwise it returns 0.

To set the SG editor to an editor with name `Ed`, call:

```
void xSGseteditor(ptr id, st Ed)
```

To display a phrase `ph`, you can call:

```
void xSGprphrase(ptr id, sgph ph)
```

Our two basic methods of interfacing to *XSG* are illustrated in short main programs in Figure 11 and Figure 12.

15 Using your own lexicons

For the format of SG lexicons, see [13]. Suppose you have such a lexicon and you want to use it with *XSG*, as an addendum to the main lexicons supplied with *XSG*. Let us be specific, and suppose that the source language is English and the name of the addendum lexicon is `ena.lx`. In general, lexicons should have file type (extension) `lx`.

The first step (after creating `ena.lx`) is to *compile* `ena.lx`. The purpose of compiling a lexicon is to produce binary files that can be loaded and accessed by *XSG* more efficiently.

If we have compiled **ESG** with the main program in either Figure 11 or Figure 12, and the name of the executable is `esg.exe`, then we can compile `ena.lx` by the command:

```
esg -compiles ena.lx
```

```
#include <xsg.h>

int main(int argc, st argv[]) {
    char seg[1000];
    sgph ph;
    ptr id;
    sint stopped = 0;

    if (xSGstart(argc, argv, &id) != 1)
        return -1;
    while (!stopped) {
        printf("Input sentence (\\"stop.\\" to stop):\n");
        gets(seg);
        if (strcmp(seg, "stop.") == 0)
            stopped = 1;
        else {
            xSGopenOut(id, "sg.out");
            ph = xSGpars(id, seg);
            fprintf(xSGout(id), "Parse of segment: \n%s\n", seg);
            xSGprphrase(id, ph);
            xSGcloseOut(id);
            system("edit sg.out");
            xSGclear(id);
        }
    }
    xSGend(id);
    return 0;
}
```

Figure 11: Interface to XSG based on xSGpars

```
#include <xsg.h>

void hook(ptr id, sgph ph) {
    fprintf(xSGout(id), "Head word of phrase = %s\n", ph->cite);
    fprintf(xSGout(id), "Showing phrase again:\n");
    xSGprphrase(id, ph);
}

int main(int argc, st argv[]) {
    ptr id;
    if (xSGstart(argc, argv, &id) != 1)
        return -1;
    xSGsetflag(id, "xout", 1);
    xSGhookfun(id, hook);
    xSGseteditor(id, "edit");
    xSGinteractive(id);
    xSGend(id);
    return 0;
}
```

Figure 12: Interface to XSG based on xSGhookfun

This will produce binary files:

```
ena.lxw, ena.lxv, ena.lxi
```

If we do not have the convenient relationship between the `main` program's command-line arguments and `xSGstart` that exists in the programs in Figure 11 and Figure 12, and instead we want to compile `ena.lx` more directly in a call to `xSGstart`, then we would set up the integer variable `argc` and the string array variable `argv`, and set

```
argc = 3;
argv[0] = "";
argv[1] = "-compiles";
argv[2] = "ena.lx";
```

and call `xSGstart(argc, argv, &id)`.

When `xSGstart` is called with the `-compiles` keyword in this way, `xSGstart` returns 0. One can see that for the programs in Figure 11 and in Figure 12, the program will not go on with the rest of its possible activities after a compile via `-compiles`. Also in this case, `xSGstart` will not try to find and load the main lexicons of *XSG*.

The compiling function `-compiles` just described does not encrypt the binary lexical files. There is another version, `-compilesenc`, which is used the same way, but applies a certain amount of encryption to the binary files.

The next step is to cause *XSG* to use `ena.lx` during parsing as an addendum to *ESG*'s main lexicon `en.lx`. Again, assuming that we have `esg.exe` compiled from the program in Figure 11 or in Figure 12, we would invoke it like so:

```
esg -mlex "ena.lx en.lx"
```

Or if we wanted to do it more directly through a call to `xSGstart`, we would set its arguments as:

```
argc = 3;
argv[0] = "";
argv[1] = "-mlex";
argv[2] = "ena.lx en.lx";
```

Although we refer here to the *source* lexicon file names (like "`ena.lx`"), the system actually reads the binary compiled forms of the lexicon.

The use of `-mlex` for specifying the lexicons in this way will cause the sequence of lexicons

```
ena.lx, en.lx
```

to be used in *preference order*. This means that when a word w is looked up during lexical analysis, **XSG** will look in order in those two lexicons (**ena.lx** first), and will use only the first successful look-up of w (ignoring later lexicons in the list for that particular w). In this sense, the addendum lexicon *overrides* the other lexicon (with this order of arguments for **-mlex**). Of course we can change what overrides what by changing the order of arguments.

There is another way to use an addendum lexicon. Instead of letting its entries override entries for the same word in the other lexicons, we can set it up so that results of look-up in the addendum lexicon are merged or “unioned” with the results (for the same word) in the other lexicons. To cause this to happen, we can call **esg** like so:

```
esg -alex ena.lx
```

Or we could make a corresponding direct call to **xSGstart**.

For an example of the use of **-alex** method, suppose we add a noun entry for the word **write** to **ena.lx**, like so:

```
write < n
```

Suppose **write** had only verb senses in the main **ESG** lexicons. Then with **ena.lx** used as an **alex**, the word **write** will get the noun sense from **ena.lx** and its verb senses from the main lexicon.

But one must be careful with **alex** lexicons. If there is too much overlap in entries with the main lexicons, then this could overload the parser with redundant analyses of words.

It is possible to specify both **alex** and **mlex** in the same call to **xSGstart**. Example for the command-line case:

```
esg -mlex "ena.lx en.lx" -alex enc.lx
```

There is another compiling function for lexicons, **-compilex**. This is simpler than **-compiles**, in that it pays no attention to multiwords in head words for entries. Actually, **-compilex** does no encryption, but there is a variant of it, **-compilexenc**, which does do encryption. If you have a source form of the ontology lexicon **ont.lx** and you need to compile that, then you must use **-compilexenc**.

This completes our discussion of user lexicons. The initialization function **xSGstart** can also invoke special utilities besides the lexical utilities discussed in this section, if suitable keywords are given to it. But these are not as important as the lexical utilities, and we will not discuss them in the current form of this report.

16 Lexical augmentation via WordNet

The lexicon `en.1x` distributed with **ESG** (in the binary files `en.1xw` and `en.1xv`) has no data derived from WordNet [14, 2]. **ESG** works well with just that base lexicon. However, it is possible to augment `en.1x` with information from WordNet, to improve performance. In this section, we describe how to do this.

WordNet is not distributed with **ESG**, but a user's installation of WordNet can be used, along with a utility `EnWNAug` packaged with the **ESG** executable `esg.exe`, to do the lexical augmentation of `en.1x` automatically. Automation of the process has the advantage that new versions of WordNet released from Princeton can be used – as long as the WordNet data files retain the same format as current ones, as in versions 2.0 and 3.0.

There are two aspects to the augmentation by `EnWNAug`: (1) Lexical entries derived from WordNet's noun vocabulary (both multiwords and single words) are added to the lexicon when they aren't in `en.1x`. (2) Semantic types, based on WordNet senses, are marked on the combined, larger lexicon.

Here's what `EnWNAug` does for (2). There is a small ontology \mathcal{O} used by **ESG**, some of whose types get marked on entries. Currently \mathcal{O} has about 130 semantic types, which I have chosen either because they are useful for **ESG** parsing, or because they are rather high-level types that might be useful to see marked in parse trees used in applications of **ESG**. Most of the types in \mathcal{O} correspond to WN (WordNet) word senses (synsets); let \mathcal{O}' be the set of those that do. The members of $\mathcal{O} \sim \mathcal{O}'$ can still be marked on parse tree nodes, because of hierarchical information in **ESG**.

The semantic types listed at the end of Section 7 are included in \mathcal{O} . They are the main ones that are useful for parsing. Some of those types are coded directly in `en.1x` and have long been there, independently of WN, although the WN augmentation sometimes results in *more* words getting marked with such types.

The basic process of type marking by `EnWNAug` on lexical entries is as follows. For each word W in the WN vocabulary, and each open-class POS (part of speech) P of W , we identify a set $S(W, P)$ of WN senses of W with POS P , as follows. Let $S'(W, P)$ be the set of *all* WN senses of W with POS P . If $S'(W, P) = \emptyset$, then we let $S(W, P) = \emptyset$. Otherwise we proceed as follows.

Basically we want to take $S(W, P)$ to be the set of all members of $S'(W, P)$ that occur with high enough (relative) *frequency*. To measure frequency, we use for any WN sense s , the corpus occurrence count $C(s)$ of s provided by WN. We let

$$T = \sum_{s \in S'(W, P)} C(s).$$

And if $T > 0$, we take the frequency of a sense $s \in S'(W, P)$ to be

$$F(s, W, P) = C(s)/T,$$

and we let

$$S(W, P) = \{s \in S'(W, P) : F(s, W, P) \geq 0.2\}.$$

And if $T = 0$, we take $S(W, P)$ to consist of just the member of $S'(W, P)$ first listed by WN. The threshold 0.2 for sense frequency is a guess at what is useful, and perhaps should be adjusted.

Now let us describe how the set of WN senses $S(W, P)$ is used to mark types from \mathcal{O}' on the augmented lexical entry for W . For each POS P of W and each $s \in S(W, P)$, we go up the WN hypernymy tree (it can branch upwards) from s , and for each WN sense s' encountered, if s' corresponds to a type $t \in \mathcal{O}'$, then we mark t on all the sense frames of W with POS P . There is a short table in **ESG** showing for each $t \in \mathcal{O}'$ what WN senses correspond to t .

Note that there is no attempt at WSD. Types appropriate to more than one sense of W (for the given POS) can get marked on W . However, the use of the frequency threshold prevents rarer senses being used. For instance, for the noun “man”, all of the senses besides the most common one (“adult male human”) have frequency numbers below the threshold. So “man” gets only the types appropriate to that most common sense. On the other hand, **ESG** parsing is *largely* syntactic (though using semi-semantic information from detailed slot frame specifications). So if “man” is used in a rare sense, the parse normally works anyway.

Now let us describe how to run **EnWNAug**. The first step is to convert the WN data files to a form that **ESG** can use. For efficiency purposes, **ESG** has its own API to WN, which I wrote. The only portion of the Princeton WN distribution needed is the set of data files, which would typically be in a directory like

```
\Program Files\WordNet\2.0\dict
```

on Windows. The **ESG** WN API needs nine data files, and expects the names to be:

```
noun.dat, noun.idx, verb.dat, verb.idx
adj.dat, adj.idx, adv.dat, adv.idx,
sense.idx
```

If your distribution of WN uses those file names, then you can use them directly where they are. But if they have another naming scheme, like

```
data.noun, index.noun, ...
```

then they should be copied to files named as above, and it is better to put these in a separate, new directory.

The next step is to execute the command

```
esg -wnmake Dir
```

from the **ESG** directory, where `Dir` is the complete path of the directory (without a final slash) where the nine WN data files are. This should create nine data files:

```
wnsg.gls, wnsg.sid, wnsg.sns, wnsg.spc,
wnsg.swd, wnsg.syn, wnsg.top, wnsg.trc, wnsg.wds
```

in the **ESG** directory. These need to be there for the **ESG** WN API to work.

The next steps are for the actual lexical augmentation. The **ESG** distribution provides the lexical files

```
en.lxw, en.lxv,
enbase.lxw, enbase.lxv
```

and these should be in the **ESG** directory. Copy (not rename) the first two files to something like

```
enSV.lxw, enSV.lxv
```

for safe keeping. It is important to save these if you wish to redo the augmentation process (e.g. with a new WN distribution), or just to run **ESG** without a WN-augmented lexicon.

Then execute:

```
esg -enwnaug
```

This should take about 15 seconds to run, and it should create binary lexical files:

```
enaug.lxw, enaug.lxv, enaug.lxi
```

This is the new, augmented form of `en.lx`. Finally, make `enaug.lx` the new form of `en.lx` by doing a copy-replace:

```
enaug.lxw => en.lxw
enaug.lxv => en.lxv
```

This completes the process of augmentation.

If you want to redo the augmentation process, then the saved base version of `en.lx`,

```
enSV.lxw, enSV.lxv,
```

should be copied (with replacement) to

```
en.lxw, en.lxv
```

and the process can be repeated.

References

- [1] Veronica Dahl and Michael C. McCord. Treating coordination in logic grammars. *Computational Linguistics*, 9:69–91, 1983.
- [2] Christiane Fellbaum. *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge, MA, 1998.
- [3] Jeffrey S. Gruber. *Studies in Lexical Relations*. PhD thesis, MIT, 1965.
- [4] Ray S. Jackendoff. *Semantics and Cognition*. MIT Press, Cambridge, MA, 1983.
- [5] Shalom Lappin and Michael C. McCord. Anaphora resolution in Slot Grammar. *Computational Linguistics*, 16:197–212, 1990.
- [6] Shalom Lappin and Michael C. McCord. A syntactic filter on pronominal anaphora for Slot Grammar. In *Proceedings of the 28th Annual Meeting of the ACL*, pages 135–142, 1990.
- [7] Michael C. McCord. Slot Grammars. *Computational Linguistics*, 6:31–43, 1980.
- [8] Michael C. McCord. Using slots and modifiers in logic grammars for natural language. *Artificial Intelligence*, 18:327–367, 1982.
- [9] Michael C. McCord. Modular logic grammars. In *Proceedings of the 23rd Annual Meeting of the ACL*, pages 104–117, 1985.
- [10] Michael C. McCord. Slot Grammar: A system for simpler construction of practical natural language grammars. In R. Studer, editor, *Natural Language and Logic: International Scientific Symposium, Lecture Notes in Computer Science*, pages 118–145. Springer Verlag, Berlin, 1990.
- [11] Michael C. McCord. Heuristics for broad-coverage natural language parsing. In *Proceedings of the ARPA Human Language Technology Workshop*, pages 127–132. Morgan-Kaufmann, 1993.
- [12] Michael C. McCord. A formal system for Slot Grammar. Technical report, IBM T. J. Watson Research Center, 2006. RC 23976.
- [13] Michael C. McCord. The Slot Grammar lexical formalism. Technical report, IBM T. J. Watson Research Center, 2006. RC 23977.
- [14] George A. Miller. WordNet: A lexical database for English. *Communications of the ACM*, 38:39–41, 1995.
- [15] Adrian Walker, Michael C. McCord, John F. Sowa, and Walter G. Wilson. *Knowledge Systems and Prolog: A Logical Approach to Expert Systems and Natural Language Processing*. Addison-Wesley, Reading, MA, 1987.