

Language Model Grammar Conversion¹

Wesley Holland, Julie Baca, Dhruva Duncan and Joseph Picone
Center for Advanced Vehicular Systems
Mississippi State University
Mississippi State MS, USA

ABSTRACT - *Supporting popular language model grammar formats, such as JSGF and XML-SRGS, has been an important step forward for the speech recognition community, particularly with respect to integration of human language technology with Internet-based technologies. Industry standard formats, though conceptually straightforward implementations of context free grammars, contain restrictions that make it nontrivial to support probabilistic finite state machines. These restrictions pose serious challenges when applied to all aspects of the speech recognition problem, such as the representation of hidden Markov models using in acoustic modeling. This paper compares and contrasts these formats, discusses the implications for speech recognition systems, and presents some solutions that have been implemented in our public domain speech recognition system.*

Keywords: {speech, grammar, SRGS, language, conversion}

1 Introduction

Several industry standard grammar specifications such as the Java Speech Grammar Format (JSGF) [1] and the W3C XML Speech Recognition Grammar Specification (XML-SRGS) [2] have been created to support development of voice-enabled Internet applications. While these standards allow for the specification of context-free grammars (CFGs), most language models for automatic speech recognition have a regular grammar equivalent and can therefore be modeled as finite state machines (FSMs). To support language model creation using these standards, we developed a suite of software tools in our public domain speech recognition toolkit [3] that convert between these grammar formats.

Issues of theoretical equivalence and restrictions on conversions between regular and context free grammars have been studied and described extensively. No algorithm

1. This material is based upon work supported by the National Science Foundation (NSF) under Grant No. IIS-0414450. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

has been proven to perform conversions from arbitrary CFGs generating regular languages to FSMs without assuming certain restrictions on the grammar, i.e. no center-embedded non-terminals [4]. However, software tools have been developed for conversions between FSMs and CFGs, which assume such restrictions on the grammars handled [5]. Nonetheless, our experience has illustrated that the specifics of individual grammar formats present unique challenges that cannot be easily addressed by theoretical solutions or generic conversions. The remainder of this paper describes the technical issues encountered in our conversion process, our solutions to these issues, and finally offers insight into future development and selection of robust, general purpose language model conversion tools.

2 Grammar Formats

As mentioned, our internal grammar format, IHD, is implemented as a set of hierarchically layered FSMs. Each FSM layer is a generic directed graph class or DiGraph. IHD binds these layers of FSMs together into levels. The top-most level contains a single DiGraph, with the nodes of this DiGraph mapping to more complete DiGraphs. Users can specify each level, but as an example, the top-most layer might represent the sentence level, the level below that the word level, the next the state level. An example is shown in Figure 1.

We first embarked on the task of providing conversion tools between the IHD format shown above and JSGF. Our goal in that task was to provide a bidirectional conversion tool for our users that could accept JSGF conformant grammars and comprehensively convert them to language models in IHD, i.e., down to the phone and state level, so

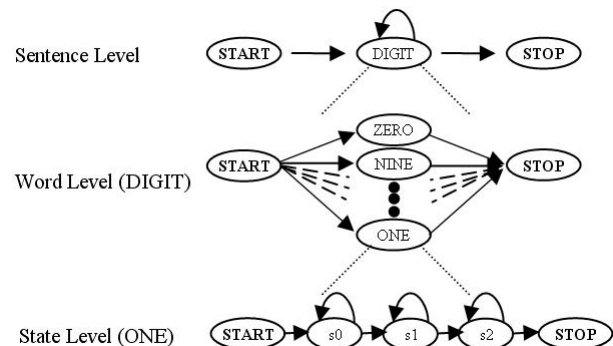


Figure 1. IHD Hierarchical Format




Figure 2. Conversion Process

that recognition experiments could be performed completely in JSGF. While not all recognition systems supporting alternate grammar formats provide this capability, we believed it was an important feature to provide our users to reduce development efforts and experimental setup time. We also required the tool to provide the same level of conversion in reverse, as shown in Figure 2. Once completed, we then undertook the task of providing a similar tool for XML

The subtle but important distinctions between JSGF and XML-SRGS and other CFG-based language models have proved challenging to the task of developing general purpose conversion tools. However, our deepened understanding of these nuances led to the design of more robust conversion tools for these and future grammar formats. The following subsection introduces key theoretical and syntactic attributes of each format, highlighting important similarities as well as differences.

2.1 BNF and EBNF

First, JSGF and XML-SRGS are theoretically equivalent in expressive and computational power, both adhering in principle to Backus-Naur Form (BNF) [6], a formal notation for CFGs, more specifically to two equivalent variants, Extended BNF (EBNF) [7] and Augmented BNF (ABNF) [7]. Many detailed descriptions of BNF and its variants exist. We provide a brief introduction to key features relevant to our discussion.

Stated simply, BNF defines a method for describing production rules in a CFG, including terminal and non-terminal symbols for rules, and a selection of alternatives among rules. Though numerous variants of the syntax exist, an example rule in a BNF grammar might be:

$$\langle A \rangle ::= \langle B \rangle | c$$

where non-terminals are represented in capital letters, A, B, surrounded by brackets $\langle \rangle$ and can appear on the lefthand side (LHS) or righthand side (RHS) of the rule demarcated by the $::=$ symbol. Terminals are often expressed in lower case (though not required), but more importantly can appear only on the rule RHS. Finally, selection or branching among alternative rule definitions is expressed by the $|$ symbol.

BNF also allows the use of recursive rules in a grammar. Such rules directly or indirectly reference themselves. An example of direct recursion might be: $\langle A \rangle ::= a \langle A \rangle$. The

use of directly or indirectly recursive rules is useful to represent repetitive actions such as loops or cycles in an FSM. Consider the simple FSM in Figure 3.

This FSM recognizes the regular expression, $a(bc)^+$ that could be represented in BNF with the production rules:

```
<S> ::= <A>
<A> ::= aB
<B> ::= bc|B
```

The use of the non-terminal B on the RHS in rule 3 is recursive and indicates that subgraph bc is a cycle that can be repeated one or more times. However, for simple regular expressions such as this, the cycle in this FSM could be represented using the $+$ operator, a standard notation for regular expressions which denotes 1 or more repetitions. EBNF extends BNF to supports the use of these structures, such as the Kleene operator $*$ and $+$ for repetition as well as others. (EBNF also has many variants, but its origins date to [7].) This allows creating a more intuitive set of production rules for regular expressions, so that rules A and B above can be reduced to:

$$\langle A \rangle ::= a (bc)^+$$

2.2 JSGF and XML-SRGS

Again, both JSGF and XML-SRGS provide expressive equivalence to EBNF. They differ, however, in syntax. As an example, the above rule could be represented in JSGF as:

$$\langle A \rangle = a(bc)^+;$$

Note that the $+$ operator is supported directly as well as the use of parentheses. Consider the same rule in XML-SRGS:

```
<rule>
  <item> a </item>
  <item repeat='1-'>
    <item> b</item>
    <item> c </item>
  </item>
</rule>
```

The $\langle rule \rangle$ tag marks this as a production rule; no non-terminal symbols are needed in this example; the terminal symbols are marked with $\langle item \rangle$ tags, and “repeat= ‘1-’” before the non-terminal b denotes the Kleene $+$ operation




Figure 3. FSM for regex $a(bc)^+$




Figure 4. FSM $(ab)^+e^+|a(ba)^*cde^+$

(1 or more repetitions) applied to the concatenation of b and c, shown by listing each non-terminal in sequence with surrounding item tags and a close item tag for repeat '1-'.
 $\langle A \rangle = a^*1(bc)$

Clearly, the JSGF syntax is more similar to EBNF than XML-SRGS and thus, familiar to designers trained in formal language theory. The differences are due in large part to their origins: XML was designed initially as a markup language for general Internet usage and later modified to provide support for spoken language; JSGF was designed from the outset to support spoken language applications. The W3C SRGS attempted to address these issues first by using JSGF as a theoretical model in defining the XML-SRGS and second, by developing a standard specification for ABNF. ABNF is an EBNF variant, with origins dating back to Arpanet. Any ABNF-SRGS can be mapped to XML-SRGS. The previous example could be written as shown below, using *1 for + before the item (bc) to be repeated:

$\langle A \rangle = a^*1(bc)$

The subtle distinctions among variants and syntax complicated the task of identifying underlying theoretical structures, but a more interesting set of issues arose with respect to the support each standard requires for recursion. While a conformant JSGF grammar processor must provide support for recursive rules, this support is optional for the XML-SRGS (and ABNF-SRGS) upon which we based our conversion. We chose, however, to support recursion for all our grammar format conversion tools. The rationale for this choice is explained further in the following subsection.

2.3 Recursion in Speech Recognition

The example in Figure 3 illustrates how a lexical construct can be represented more simply with a regular expression than a recursive CFG. However, the converse situation often arises in speech recognition: though a lexical or syntactic construct could be described by a complex regular expression, the recursive form may be simpler to produce algorithmically. Figure 4 shows such an example.

Though still a relatively simple graph, notice the branches, cycles within branches, and self-loops not present in Figure 3. The regular expression for the strings accepted by this FSA can be written as: $(ab)^+e^+|a(ba)^*cde^+$. A rule for this can be written in EBNF, without non-terminals and without recursion, as follows:

$S ::= (ab)^+e^+ | a(ba)^*cde^+$

Production rules can also be written for the expression using non-terminals, direct recursion for the self-loop, and indirect recursion for the cycle by creating a non-terminal for each state in the graph, e.g., A for a, as follows:

$S ::= A$
 $A ::= a(B|C)$
 $B ::= b(A|E)$
 $C ::= cD$
 $D ::= dE$
 $E ::= eE | ET$
 $T ::= t$

While the single rule in the first grammar eases visualizing a legal input string directly from the rules, the second grammar more directly expresses actions to be encoded in an algorithm with limited lookahead, e.g., rule A states that on input 'a', branch to either state B or C, while the indirectly recursive rule B states on input 'b', branch back to A or ahead to E.

The use of recursive rules is an attribute that distinguishes regular and context free grammars. As noted, although XML and JSGF both provide the expressive power of CFGs, only JSGF requires support for recursion. The restrictions on the type of recursion supported should be noted, however. Notice that the recursive grammar for Figure 4 is right recursive, that is, all recursive references appear on the rightmost side of the RHS. JSGF limits its required support to these type grammars. Several arguments can be made in favor of this restriction. First, any right recursive rule can be rewritten using the Kleene * and + operators where more appropriate. Second, speech recognizers typically use regular grammars, which must be either left or right linear, and thus can contain only left or right recursion. Finally, right recursive grammars can be parsed by top-down parsers with limited lookahead, such as LL(1), which are arguably among the simplest to construct.

A final example of recursion taken from our speech recognition system further explains our decision to implement support for both the EBNF grammar structures, e.g., * and +, as well as right recursive rules. To restate, we sought to provide conversion tools that would allow use of XML-SRGS and JSGF (and other equivalent grammar formats) to perform complete recognition experiments, i.e., all the way from the sentence level down to the acoustic and phone level. This meant, however, writing and parsing grammars that represented layers of FSMs, with those at




Figure 5. FSM for digits grammar




Figure 6. FSM for XML weights and probabilities

the state level containing many self-loops. Consider a sentence level view of a simple digits grammar in Figure 5.

Assuming SILEN and DIGIT are terminal symbols, a simple regular expression (SILEN DIGIT*)+ can describe this FSM. However, SILEN and DIGIT are actually non-terminals describing inputs that are modeled at lower levels. DIGIT is represented at the word and state levels as shown in Figure 1. It is possible to write production rules containing regular expressions for each of the sentence, word and state levels. Our system, however, treats non-terminals on arcs as subroutines, saving the current location in the graph, processing the non-terminals, and then returning when parsed, behaving effectively as recursive transition networks (RTNs). This meant that our conversion tool for IHD->XML could easily generate grammars with either type of structure, e.g., *, + for self-loops and cycles, or recursive rules. RTNs are commonly used in spoken language processing to represent language models that are predominantly FSM-based with some recursion. This further reinforced our decision to support both recursive rules and EBNF extensions.

3 Conversion Redesign

Other specific features of XML complicated the design decisions described in the previous section. As shown, the syntax for loops and cycles written EBNF-style differed from that using recursive rules. More importantly, the handling of weights on arcs differed subtly. Since weights are a critical component of ASR systems, this was an important issue.

3.1 XML-SRGS Weights and Probabilities

Two methods are available for specifying weights in XML-SRGS: the weight attribute and the repeat-prob attribute. It is important to note that a repeat probability is a different logical entity from a weight. It is the probability that a given loop will repeat, while a weight can only be transformed into a probability when compared with the other weights leaving a node. The FSM in

Figure 6 is used to illustrate both methods. An XML-SRGS grammar using both the weight attribute and the repeat-prob attribute is shown.

```

<grammar>
  <rule>
    <item> a </item>
    <one-of>
      <item weight='4'> b </item>
  
```

```

    </one-of>
    <item repeat='1-' repeat_prob='.43'> c </item>
  </rule>
</grammar>

```

The first item in this grammar, a, has no weight since its incoming arc has no weight. The second item, b, has a weight on its incoming arc. SRGS dictates that an item may not have a weight unless its immediate enclosing tag is a <one-of>, the rationale being that weights are not needed unless alternatives are present. SRGS does, however, allow for a single item to be enclosed in a <one-of> in order to specify a weight where one would not otherwise be allowed. Since speech recognition systems typically have weights on all arcs, this limitation on the SRGS weight attribute is significant.

Also important, the repeat-prob attribute can only be used with the the SRGS repeat looping attribute. Recall this attribute implements the EBNF loop extensions for Kleene operations. In this example, “repeat = ‘1-’” is equivalent to +. An additional limitation is that repeat probability values must lie between 0.0 and 1.0.

The above grammar would be represented as shown below without repeat probabilities, but with recursive rules:

```

<grammar>
  <rule>
    <item> a </item>
    <one-of>
      <item weight='4'> <ruleref uri="#B"/> </item>
    </one-of>
  </rule>
  <rule id="B">
    <item> b </item>
    <item> <ruleref uri="#C"/> </item>
  </rule>
  <rule id="C">
    <item> c </item>
    <one-of>
      <item weight='.43'> <ruleref uri="#C"/> </item>
      <item> <ruleref special="NULL"/> </item>
    </one-of>
  </rule>
</grammar>

```

This avoids the use of a repeat probability on node ‘c’ by putting a weight on the recursive rule reference at the end of this grammar. However, special care must be taken in converting weight and repeat-prob attributes consistently. For example, our system uses only log probabilities at the state level. Converting the original grammar to represent all weights consistently was critical. This would require processing the grammar to represent weights in a uniform manner.

3.2 ABNF and BNF Conversion Modules

Once the underlying theoretical structures of each format were understood in detail, it was apparent that producing a verifiably robust conversion tool would require additional stages, and hence conversion modules, in our process. The first stage would create a common EBNF grammar format to which any other format could be converted for verification of correctness. The ABNF-SRGS was an obvious choice as the common EBNF format. The next stage would entail processing the EBNF to remove extensions and thus standardize representation of weights and probabilities. We redesigned our conversion process to include the following steps and corresponding software modules: 1) convert the XML to an equivalent ABNF, 2) convert the ABNF to remove the EBNF extensions and produce a clean BNF with or without recursion, 3) convert the clean BNF to IHD. We also created modules to implement the steps shown above in the reverse conversion from IHD \rightarrow XML. The redesigned conversion process is shown in Figure 7. We are redesigning our JSGF conversion tools to include these stages and corresponding modules as well.

4 Future Work

Our redesigned conversion suite will not only extend the language model compatibility of our system, but also allows for easy conversion between popular grammar formats using BNF as an intermediary. Use of this theoretically grounded intermediary will allow for easy application of textbook algorithms to modern industrial grammar formats. In addition to these overall benefits, each supported format will bring with it individual benefits. Support for ABNF will mark the realization of a theoretical format as a practical recognition grammar. Support for JSGF will provide an easy interface with current generation speech recognizers. Lastly, XML-SRGS support will pave the way for future web applications of our speech recognition software.

5 Summary and Conclusions

Supporting popular CFG-based language model formats has been an important priority in our research. When we undertook this task, we anticipated the theoretical conversion between CFG and regular grammar formats to be our greatest technical challenge. Instead, the important nuances of each CFG format implementation have presented equal challenges to producing robust conversion tools. Through a comprehensive approach of supporting conversion down to the phone and state level, we have gained a detailed understanding of the nuances among formats. These nuances added significant complexity to the conversion process, though in theory, all CFG formats offer equivalent expressive power. In particular, the implementation differences with respect to EBNF and BNF grammar styles and support for recursion are critical.




Figure 7. Conversion redesign

Further, assignment of weights and probabilities for these styles must be carefully considered in converting to any other formats, including CFG or regular grammars. We have addressed these issues by incorporating additional stages and corresponding software modules in our process which perform generic conversions to and from common EBNF and BNF grammar formats. These enhancements and the knowledge gained from our work have advanced an important goal for the speech research community—producing verifiably robust conversion tools to support current and future popular CFG-based language model standards.

6 References

- [1] Java Speech Grammar Format Specification, Version 1, Sun Microsystems Developer Network, October 26, 1998 (see <http://java.sun.com/products/java-media/speech/forDevelopers/JSGF/JSGF>).
- [2] A. Hunt and S. McGlashan, Eds., W3C Speech Recognition Grammar Specification Version 1.0, March 16, 2004 (see <http://www.w3.org/TR/speech-grammar/>).
- [3] J. Picone, et al., “A Public Domain C++ Speech Recognition Toolkit,” ISIP, Mississippi State University, Mississippi State, MS, USA, March 2003 (<http://www.isip.msstate.edu/projects/speech>).
- [4] N. Chomsky, “On Certain Formal Properties of Grammars,” *Information and Control*, Vol. 2, 1959, pp. 137-167.
- [5] Mohri, M, “Weighted Grammar Tools: The GRM Library,” in J.C. Junqua and G. van Noord (eds), *Robustness in Language and Speech Technology*, Kluwer Academic Publishers.
- [6] Naur, P. “Revised Report on the Algorithmic Language Algol 60,” *Communications of the ACM*, Vol. 7, No. 12, 1963, pp. 735–736.
- [7] Wirth, N. “What Can We Do About the Unnecessary Diversity of Notation for Syntactic Definitions,” *Communications of the ACM*, Vol. 20, No. 11, 1977, pp. 822–823.