

Tabular Parsing

Mark-Jan Nederhof*

Faculty of Arts

University of Groningen

P.O. Box 716

NL-9700 AS Groningen, The Netherlands

`markjan@let.rug.nl`

Giorgio Satta

Department of Information Engineering

University of Padua

via Gradenigo, 6/A

I-35131 Padova, Italy

`satta@dei.unipd.it`

1 Introduction

Parsing is the process of determining the parses of an input string according to a grammar. In this chapter we will restrict ourselves to context-free grammars. Parsing is related to *recognition*, which is the process of determining whether an input string is in the language described by a grammar or automaton. Most algorithms we will discuss are recognition algorithms, but since they can be straightforwardly extended to perform parsing, we will not make a sharp distinction here between parsing and recognition algorithms.

For a given grammar and an input string, there may be very many parses,

*Supported by the Royal Netherlands Academy of Arts and Sciences. Secondary affiliation is the German Research Center for Artificial Intelligence (DFKI).

perhaps too many to be enumerated one by one. Significant practical difficulties in computing and storing the parses can be avoided by computing individual fragments of these parses and storing them in a table. The advantage of this is that one such fragment may be shared by many different parses. The methods of *tabular parsing* that we will investigate in this chapter are capable of computing and representing exponentially many parses in polynomial time and space, respectively, by means of this idea of sharing of fragments between several parses.

Tabular parsing, invented in the field of computer science in the period roughly between 1965 and 1975, also became known later in the field of computational linguistics as *chart parsing* [36]. Tabular parsing is a form of *dynamic programming*. A very related approach is to apply *memoization* to functional parsing algorithms [21].

What is often overlooked in modern parsing literature is that many techniques of tabular parsing can be straightforwardly derived from non-tabular parsing techniques expressed by means of push-down automata. A push-down automaton is a device that reads input from left to right, while manipulating a stack. Stacks are a very common data structure, frequently used wherever there is recursion, such as for the implementation of functions and procedures in programming languages, but also for context-free parsing.

Taking push-down automata as our starting point has several advantages for describing tabular parsers. Push-down automata are simpler devices than the tabular parsers that can be derived from them. This allows us to get acquainted with simple, non-tabular forms of context-free parsing before we move on to tabulation, which can, to a large extent, be explained independently from the workings of individual push-down automata. Thereby we achieve a separation of concerns. Apart from these presentational advantages, parsers can also be implemented more easily with this modular design than without.

In Section 2 we discuss push-down automata and their relation to context-free grammars. Tabulation in general is introduced in Section 3. We then discuss a small number of specific tabular parsing algorithms that are well-known in the literature, viz. Earley’s algorithm (Section 4), the Cocke-Kasami-Younger algorithm (Section 5), and tabular LR parsing (Section 6). Section 7 discusses compact representations of sets of parse trees, which can be computed by tabular parsing algorithms. Section 8 provides further point-

ers to relevant literature.

2 Push-down automata

The notion of push-down automaton plays a central role in this chapter. Contrary to what we find in some textbooks, our push-down automata do not possess states next to stack symbols. This is without loss of generality, since states can be encoded into the stack symbols. Thus, a *push-down automaton* (PDA) $\mathcal{A}$ is a 5-tuple $(\Sigma, Q, q_{init}, q_{final}, \Delta)$, where Σ is an alphabet, i.e., a finite set of input symbols, Q is a finite set of *stack symbols*, including the *initial stack symbol* q_{init} and the *final stack symbol* q_{final} , and Δ is a finite set of *transitions*.

A transition has the form $\sigma_1 \xrightarrow{v} \sigma_2$, where $\sigma_1, \sigma_2 \in Q^*$ and $v \in \Sigma^*$. Such a transition can be applied if the stack symbols σ_1 are found to be the top-most few symbols on the stack and the input symbols v are the first few symbols of the unread part of the input. After application of such a transition, σ_1 has been replaced by σ_2 , and the next $|v|$ input symbols are henceforth treated as having been read.

More precisely, for a fixed PDA and a fixed input string $w = a_1 \cdots a_n \in \Sigma^*$, $n \geq 0$, we define a *configuration* as a pair (σ, i) consisting of a stack $\sigma \in Q^*$ and an input position i , $0 \leq i \leq n$. The *input position* indicates how many of the symbols from the input have already been read. Thereby, position 0 and position n indicate the beginning and the end, respectively, of w . We define the binary relation $\vdash$ on configurations by: $(\sigma, i) \vdash (\sigma', j)$ if and only if there is some transition $\sigma_1 \xrightarrow{v} \sigma_2$ such that $\sigma = \sigma_3 \sigma_1$ and $\sigma' = \sigma_3 \sigma_2$, some $\sigma_3 \in Q^*$, and $v = a_{i+1} a_{i+2} \cdots a_j$. Here we assume $i \leq j$, and if $i = j$ then $v = \varepsilon$, where ε denotes the empty string. Note that in our notation, stacks grow from left to right, i.e., the top-most stack symbol will be found at the right end.

We denote the reflexive and transitive closure of $\vdash$ by $\vdash^*$; in other words, $(\sigma, i) \vdash^* (\sigma', j)$ means that we may obtain configuration (σ', j) from (σ, i) by applying zero or more transitions. We say that the PDA *recognizes* a string $w = a_1 \cdots a_n$ if $(q_{init}, 0) \vdash^* (q_{final}, n)$. This means that we start with a stack containing only the initial stack symbol, and the input position is initially 0, and recognition is achieved if we succeed in reading the complete input, up to the last position n , while the stack contains only the final stack symbol.

$$\begin{array}{llll}
q_0 \xrightarrow{a} q_0 q_1 & q_2 \xrightarrow{c} q_2 q_4 & q_4 q_5 \xrightarrow{\varepsilon} q_6 & q_3 q_6 \xrightarrow{\varepsilon} q_8 \\
q_0 q_1 \xrightarrow{b} q_0 q_2 & q_3 \xrightarrow{c} q_3 q_4 & q_2 q_6 \xrightarrow{\varepsilon} q_7 & q_0 q_8 \xrightarrow{\varepsilon} q_9 \\
q_0 q_1 \xrightarrow{b} q_0 q_3 & q_4 \xrightarrow{d} q_4 q_5 & q_0 q_7 \xrightarrow{\varepsilon} q_9 &
\end{array}$$

Figure 1: Transitions of an example PDA.

q_0	0	q_0	0
$q_0 q_1$	1	$q_0 q_1$	1
$q_0 q_2$	2	$q_0 q_3$	2
$q_0 q_2 q_4$	3	$q_0 q_3 q_4$	3
$q_0 q_2 q_4 q_5$	4	$q_0 q_3 q_4 q_5$	4
$q_0 q_2 q_6$	4	$q_0 q_3 q_6$	4
$q_0 q_7$	4	$q_0 q_8$	4
q_9	4	q_9	4

Figure 2: Two sequences of configurations, leading to recognition of the string $abcd$.

The language *accepted* by a PDA is the set of all strings that it recognizes.

As an example, consider the PDA with $\Sigma = \{a, b, c, d\}$, $Q = \{q_0, \dots, q_9\}$, $q_{init} = q_0$, $q_{final} = q_9$, and the set Δ of transitions given in Figure 1. There are two ways of recognizing the input string $w = a_1 a_2 a_3 a_4 = abcd$, indicated by the two sequences of configurations in Figure 2.

We say a PDA is *deterministic* if for each configuration there can be at most one applicable transition. The example PDA above is clearly *non-deterministic* due to the two transitions $q_0 q_1 \xrightarrow{b} q_0 q_2$ and $q_0 q_1 \xrightarrow{b} q_0 q_3$.

A *context-free grammar* (CFG) $\mathcal{G}$ is a 4-tuple (Σ, N, S, R) , where Σ is an *alphabet*, i.e., a finite set of *terminals*, N is a finite set of *nonterminals*, including the *start symbol* S , and R is a finite set of *rules*, each of the form $A \rightarrow \alpha$ with $A \in N$ and $\alpha \in (\Sigma \cup N)^*$. The usual ‘derives’ relation is denoted by $\Rightarrow$, and its reflexive and transitive closure by $\Rightarrow^*$. The language *generated* by a CFG is the set $\{w \mid S \Rightarrow^* w\}$.

In practice, a PDA is not hand-written, but is automatically obtained from a CFG, by a mapping that preserves the generated/accepted language. Particular mappings from CFGs to PDAs can be seen as formalizations of

parsing strategies.

We define the size of a PDA as $\sum_{(\sigma_1 \xrightarrow{v} \sigma_2) \in \Delta} |\sigma_1 v \sigma_2|$, i.e., the total number of occurrences of stack symbols and input symbols in the set of transitions. Similarly, we define the size of a CFG as $\sum_{(A \rightarrow \alpha) \in R} |A \alpha|$, i.e., the total number of occurrences of terminals and nonterminals in the set of rules.

3 Tabulation

In this section, we will restrict the allowable transitions to those of the types $q_1 \xrightarrow{a} q_1 q_2$, $q_1 q_2 \xrightarrow{a} q_1 q_3$, and $q_1 q_2 \xrightarrow{\varepsilon} q_3$, where $q_1, q_2, q_3 \in Q$ and $a \in \Sigma$. The reason is that this allows a very simple form of tabulation, based on work by [17, 6]. In later sections, we will again consider less restrictive types of transitions. Note that each of the transitions in Figure 1 is of one of the three types above.

The two sequences of configurations in Figure 2 share a common step, viz. the application of transition $q_4 \xrightarrow{d} q_4 q_5$ at input position 3 when the top-of-stack is q_4 . In this section we will show how we can avoid doing this step twice. Although the savings in time and space for this toy example are negligible, in realistic examples we can reduce the costs from exponential to polynomial, as we will see later.

A central observation is that if two configurations share the same top-of-stack and the same input position, then the sequences of steps we can perform on them are identical as long as we do not access lower regions of the stack that differ between these two configurations. This implies for example that in order to determine which transition(s) of the form $q_1 \xrightarrow{a} q_1 q_2$ to apply, we only need to know the top-of-stack q_1 , and the current input position so that we can check whether a is the next unread symbol from the input.

These considerations lead us to propose a representation of sets of configurations as graphs. The set of vertices is partitioned into subsets, one for each input position, and each such subset contains at most one vertex for each stack symbol. This last condition is what will allow us to share steps between different configurations.

We also need arcs in the graph to connect the stack symbols. This is necessary when transitions of the form $q_1 q_2 \xrightarrow{a} q_1 q_3$ or $q_1 q_2 \xrightarrow{\varepsilon} q_3$ are applied, since these require access to deeper regions of a stack than just its




Figure 3: The collection of all derivable configurations represented as graph. For each input position there is a subset of vertices. For each such subset, there is at most one vertex for each stack symbol.

top symbol. The graph will contain an arc from a vertex representing stack symbol q at position i to a vertex representing stack symbol q' at position $j \leq i$, if q' resulted as the topmost stack symbol at input position j , and q can be immediately on top of that q' at position i . If we take a path from a vertex in the subset of vertices for position i , and follow arcs until we cannot go any further, encountering stack symbols $q_1, \dots, q_m$, in this order, then this means that $(q_{init}, 0) \vdash^* (q_m \cdots q_1, i)$.

For the running example, the graph after completion of the parsing process is given in Figure 3. One detail we have not yet mentioned is that we need an imaginary stack symbol $\perp$, which we assume occurs below the actual bottom-of-stack. We need this symbol to represent stacks consisting of a single symbol. Note that the path from the vertex labelled q_9 in the subset for position 4 to the vertex labelled $\perp$ means that $(q_0, 0) \vdash^* (q_9, 4)$, which implies the input is recognized.

What we still need to explain is how we can construct the graph, for a given PDA and input string. Let $w = a_1 \cdots a_n$, $n \geq 0$, be an input string. In the algorithm that follows, we will manipulate 4-tuples (j, q', i, q) , where

1. Let $\mathcal{T} = \{(\perp, 0, q_{init}, 0)\}$.
2. For $i = 1, \dots, n$ do:
 - (a) Let $\mathcal{N} = \emptyset$.
 - (b) For each $(q', j, q_1, i-1) \in \mathcal{T}$ and each transition $q_1 \xrightarrow{a_i} q_2$ such that $(q_1, i-1, q_2, i) \notin \mathcal{T}$, add $(q_1, i-1, q_2, i)$ to $\mathcal{T}$ and to $\mathcal{N}$.
 - (c) For each $(q_1, j, q_2, i-1) \in \mathcal{T}$ and each transition $q_1 q_2 \xrightarrow{a_i} q_1 q_3$ such that $(q_1, j, q_3, i) \notin \mathcal{T}$, add (q_1, j, q_3, i) to $\mathcal{T}$ and to $\mathcal{N}$.
 - (d) As long as $\mathcal{N} \neq \emptyset$ do:
 - i. Remove some (q_1, j, q_2, i) from $\mathcal{N}$.
 - ii. For each $(q', k, q_1, j) \in \mathcal{T}$ and each transition $q_1 q_2 \xrightarrow{\varepsilon} q_3$ and $(q', k, q_3, i) \notin \mathcal{T}$, add (q', k, q_3, i) to $\mathcal{T}$ and to $\mathcal{N}$.
3. Recognize the input if $(\perp, 0, q_{final}, n) \in \mathcal{T}$.

Figure 4: Tabular algorithm to find the collection of all derivable configurations for input $a_1 \dots a_n$, in the form of a set $\mathcal{T}$ of items.

$q', q \in Q$ and j, i are input positions with $0 \leq j \leq i \leq n$. These 4-tuples will be called *items*. Item (j, q', i, q) means that there is an arc in the graph from a vertex representing q at position i to a vertex representing q' at position j . Formally, it means that for some σ we have $(q_{init}, 0) \vdash^* (\sigma \ q', j)$ and $(\sigma \ q', j) \vdash^* (\sigma \ q' \ q, i)$, where in the latter relation the transitions that are involved do not access any symbols internal to σ .

The algorithm is given in Figure 4. Initially, we let the set $\mathcal{T}$ contain only the item $(\perp, 0, q_{init}, 0)$, representing one arc in the graph. We then incrementally fill $\mathcal{T}$ with more items, representing more arcs in the graph, until the complete graph has been constructed. In this particular tabular algorithm, we process the symbols from the input one by one, from left to right, applying all transitions as far as we can before moving on to the next input symbol. Whereas $\mathcal{T}$ contains all items that have been derived up to a certain point, the set $\mathcal{N}$ contains only those items from $\mathcal{T}$ that still need to be combined with others in order to (possibly) obtain new items. The set $\mathcal{T}$ will henceforth be called the *table* and the set $\mathcal{N}$ the *agenda*.

Let us analyze the worst-case time complexity of the algorithm in Fig-

ure 4. We assume that the table $\mathcal{T}$ is implemented as a square array of size $n + 1$, indexed by input positions i and j , and that each item can be stored in and retrieved from $\mathcal{T}$ in time $\mathcal{O}(1)$. The agenda $\mathcal{N}$ can be implemented as a stack. Let us consider Step 2(d). A single application of this step takes time $\mathcal{O}(1)$. Since each such application is uniquely identified by a transition $q_1 q_2 \xrightarrow{\varepsilon} q_3$, a stack symbol q' and the three input positions i, j and k , the number of possible applications of the step is $\mathcal{O}(|\Delta| |Q| n^3)$, which for our PDAs can be rewritten as $\mathcal{O}(|\mathcal{A}| |Q| n^3)$. It is not difficult to see that this quantity also dominates the worst-case time complexity of our algorithm, which is thereby polynomial both in the size of the PDA and in the length of the input string. A similar analysis shows that the space complexity of the algorithm is $\mathcal{O}(|Q|^2 n^2)$.

Although the use of the agenda in the algorithm from Figure 4 allows a fairly straightforward implementation, it obscures somewhat how items are derived from other items. This can be described more clearly by abstracting away from certain details of the algorithm, such as the order in which items are added to $\mathcal{T}$. This can be achieved by means of a *deduction system* [31].¹ Such a system contains a set of inference rules, each consisting of a list of *antecedents*, which stand for items that we have already established to be in $\mathcal{T}$, and, below a horizontal line, the *consequent*, which stands for an item that we derive from the antecedents and that is added to $\mathcal{T}$ unless it is already present. At the right of an inference rule, we may also write a number of *side conditions*, which indicate when rules may be applied, on the basis of transitions of the PDA.

A deduction system equivalent to the algorithm from Figure 4 is given in Figure 5. In Figure 3, $(0, q_0, 1, q_1)$ is derived from $(\perp, 0, q_0, 0)$ by means of $q_0 \xrightarrow{a} q_0 q_1$, a being a_1 ; $(0, q_0, 2, q_2)$ is derived from $(0, q_0, 1, q_1)$ by means of $q_0 q_1 \xrightarrow{b} q_0 q_2$, b being a_2 ; $(0, q_0, 4, q_7)$ is derived from $(0, q_0, 2, q_2)$ and $(2, q_2, 4, q_6)$ by means of $q_2 q_6 \xrightarrow{\varepsilon} q_7$.

We may now extend our repertoire of transitions by those of the forms $q_1 \xrightarrow{\varepsilon} q_1 q_2$ and $q_1 q_2 \xrightarrow{\varepsilon} q_1 q_3$, which only requires two additional inference rules, indicated in Figure 6. To extend the algorithm in Figure 4 to handle these additional types of transitions requires more effort. Up to now, all items (q, j, q', i) , with the exception of $(\perp, 0, q_{init}, 0)$, were such that $j < i$.

¹The earliest mention of abstract specifications of parsing algorithms may be due to [8]. See also [32].

$$\begin{array}{c}
\frac{}{(\perp, 0, q_{init}, 0)} \qquad \frac{(q_1, j, q_2, i-1)}{(q_1, j, q_3, i)} \left\{ \begin{array}{l} q_1 \ q_2 \xrightarrow{a_i} q_1 \ q_3 \end{array} \right. \\
\\
\frac{(q', j, q_1, i-1)}{(q_1, i-1, q_2, i)} \left\{ \begin{array}{l} q_1 \xrightarrow{a_i} q_1 \ q_2 \end{array} \right. \qquad \frac{(q', k, q_1, j)}{(q_1, j, q_2, i)} \left\{ \begin{array}{l} q_1 \ q_2 \xrightarrow{\varepsilon} q_3 \end{array} \right.
\end{array}$$

Figure 5: Tabular parsing algorithm in the form of a deduction system.

$$\frac{(q', j, q_1, i)}{(q_1, i, q_2, i)} \left\{ \begin{array}{l} q_1 \xrightarrow{\varepsilon} q_1 \ q_2 \end{array} \right. \qquad \frac{(q_1, j, q_2, i)}{(q_1, j, q_3, i)} \left\{ \begin{array}{l} q_1 \ q_2 \xrightarrow{\varepsilon} q_1 \ q_3 \end{array} \right.$$

Figure 6: Two additional inference rules for transitions of the form $q_1 \xrightarrow{\varepsilon} q_1 \ q_2$ and $q_1 \ q_2 \xrightarrow{\varepsilon} q_1 \ q_3$.

If we had an item (q_1, j, q_2, i) in the agenda $\mathcal{N}$ and were looking for items (q', k, q_1, j) in $\mathcal{T}$, in order to apply a transition $q_1 \ q_2 \xrightarrow{\varepsilon} q_3$, then we could be sure that we had access to all items (q', k, q_1, j) that would ever be added to $\mathcal{T}$. This is because $j < i$, and all items having j as second input position had been found at an earlier iteration of the algorithm.

However, if we add transitions of the form $q_1 \xrightarrow{\varepsilon} q_1 \ q_2$ and $q_1 \ q_2 \xrightarrow{\varepsilon} q_1 \ q_3$, we may obtain items of the form (q, j, q', i) with $j = i$. It may then happen that an item (q', k, q_1, j) is added to $\mathcal{T}$ *after* the item (q_1, j, q_2, i) is taken from the agenda $\mathcal{N}$ and processed. To avoid that we overlook any computation of the PDA, we must change the algorithm to take into account that an item taken from the agenda may be of the form (q', k, q_1, j) , and we then need to find items of the form (q_1, j, q_2, i) already in the table, with $j = i$, in order to apply a transition $q_1 \ q_2 \xrightarrow{\varepsilon} q_3$. We leave it to the reader to determine the precise changes this requires to Figure 4, and to verify that it is possible to implement these changes in such a way that the order of the time and space complexity remains unchanged.

4 Earley's algorithm

In this section we will investigate the top-down parsing strategy, and discuss tabulation of the resulting PDAs. Let us fix a CFG $\mathcal{G} = (\Sigma, N, S, R)$ and let us assume that there is only one rule in R of the form $S \rightarrow \alpha$. The stack

symbols of the PDA that we will construct are the so called *dotted rules*, defined as symbols of the form $A \rightarrow \alpha \bullet \beta$ where $A \rightarrow \alpha\beta$ is a rule from R ; in words, a stack symbol is a rule in which a dot has been inserted somewhere in the right-hand side. Intuitively, the dot separates the grammar symbols that have already been found to derive substrings of the read input from those that are still to be processed. We will sometimes enclose dotted rules in round brackets to enhance readability.

The alphabet of the PDA is the same as that of the CFG. The initial stack symbol is $S \rightarrow \bullet \alpha$, the final stack symbol is $S \rightarrow \alpha \bullet$, and the transitions are:

1. $(A \rightarrow \alpha \bullet B\beta) \xrightarrow{\varepsilon} (A \rightarrow \alpha \bullet B\beta) (B \rightarrow \bullet \gamma)$ for all rules $A \rightarrow \alpha B\beta$ and $B \rightarrow \gamma$;
2. $(A \rightarrow \alpha \bullet b\beta) \xrightarrow{b} (A \rightarrow \alpha b \bullet \beta)$ for each rule $A \rightarrow \alpha b\beta$, where $b \in \Sigma$;
3. $(A \rightarrow \alpha \bullet B\beta) (B \rightarrow \gamma \bullet) \xrightarrow{\varepsilon} (A \rightarrow \alpha B \bullet \beta)$ for all rules $A \rightarrow \alpha B\beta$ and $B \rightarrow \gamma$.

Given a stack symbol $A \rightarrow \alpha \bullet X\beta$, with $X \in \Sigma \cup N$, the indicated occurrence of X will here be called the *goal*. The goal in the top-of-stack is the symbol that must be matched against the next few unread input symbols. Transitions of type 1 above predict rules with nonterminal B in the left-hand side, when B is the goal in the top-of-stack. Transitions of type 2 move the dot over terminal goal b in the top-of-stack, if that b matches the next unread input symbol. Finally, transitions of type 3 combine the top-most two stack symbols, when the top-of-stack indicates that the analysis of a rule with B in the left-hand side has been completed. The current top-of-stack is removed, and in the new top-of-stack, the dot is moved over the goal B .

Since the types of transition above are covered by what we discussed in Section 3, we may apply a subset of the inference rules from Figures 5 and 6 to obtain a tabular parsing algorithm for the top-down strategy. This will result in items of the form

$$(A \rightarrow \alpha \bullet B\beta, j, B \rightarrow \gamma \bullet \delta, i).$$

However, it can be easily verified that if there is such an item in the table, and if some stack symbol $A' \rightarrow \alpha' \bullet B\beta'$ may occur on top of the stack at

$$\frac{}{(0, S \rightarrow \bullet \alpha, 0)} \left\{ S \rightarrow \alpha \right. \quad (1)$$

$$\frac{(j, A \rightarrow \alpha \bullet B\beta, i)}{(i, B \rightarrow \bullet \gamma, i)} \left\{ B \rightarrow \gamma \right. \quad (2)$$

$$\frac{(j, A \rightarrow \alpha \bullet b\beta, i-1)}{(j, A \rightarrow \alpha b \bullet \beta, i)} \left\{ b = a_i \right. \quad (3)$$

$$\frac{\begin{array}{c} (k, A \rightarrow \alpha \bullet B\beta, j) \\ (j, B \rightarrow \gamma \bullet, i) \end{array}}{(k, A \rightarrow \alpha B \bullet \beta, i)} \quad (4)$$

Figure 7: Tabular top-down parsing, or Earley's algorithm.

position j , then at some point, the table will also contain the item

$$(A' \rightarrow \alpha' \bullet B\beta', j, B \rightarrow \gamma \bullet \delta, i).$$

An implication of this is that the first component $A \rightarrow \alpha \bullet B\beta$ of an item represents redundant information, and may be removed without affecting the correctness of the tabular algorithm. (See [27, Section 1.2.2] for the exact conditions that justify this simplification.) After this simplification, we obtain the deduction system in Figure 7, which can be seen as a specialized form of the tabular algorithm from the previous section. It is also known as Earley's algorithm [9, 2, 11]. Step (1) is called *initializer*, (2) is called *predictor*, (3) is called *scanner*, and (4) is called *completer*.

As an example, consider the CFG with $\Sigma = \{a, *, +\}$, $N = \{S, E\}$ and with rules $S \rightarrow E$, $E \rightarrow E * E$, $E \rightarrow E + E$ and $E \rightarrow a$, and consider the input string $w = a + a * a$. Now that items are 3-tuples, it is more convenient to represent the table $\mathcal{T}$ as an upper triangular matrix rather than a graph, as exemplified by Figure 8. This matrix consists of sets $\mathcal{T}_{i,j}$, $i \leq j$, such that $(A \rightarrow \alpha \bullet \beta) \in \mathcal{T}_{i,j}$ if and only if $(i, A \rightarrow \alpha \bullet \beta, j) \in \mathcal{T}$.

	0	1	2	3	4	5
0	$S \rightarrow \bullet E$ $E \rightarrow \bullet E * E$ $E \rightarrow \bullet E + E$ $E \rightarrow \bullet a$	$E \rightarrow a \bullet$ $S \rightarrow E \bullet$ $E \rightarrow E \bullet * E$ $E \rightarrow E \bullet + E$	$E \rightarrow E + \bullet E$	$E \rightarrow E + E \bullet$ $S \rightarrow E \bullet$ $E \rightarrow E \bullet * E$ $E \rightarrow E \bullet + E$	$E \rightarrow E * \bullet E$	$E \rightarrow E * E \bullet$ $E \rightarrow E + E \bullet$ $S \rightarrow E \bullet$ $E \rightarrow E \bullet * E$ $E \rightarrow E \bullet + E$
1						
2			$E \rightarrow \bullet E * E$ $E \rightarrow \bullet E + E$ $E \rightarrow \bullet a$	$E \rightarrow a \bullet$ $S \rightarrow E \bullet$ $E \rightarrow E \bullet * E$ $E \rightarrow E \bullet + E$	$E \rightarrow E * \bullet E$	$E \rightarrow E * E \bullet$ $E \rightarrow E \bullet * E$ $E \rightarrow E \bullet + E$
3						
4					$E \rightarrow \bullet E * E$ $E \rightarrow \bullet E + E$ $E \rightarrow \bullet a$	$E \rightarrow a \bullet$ $E \rightarrow E \bullet * E$ $E \rightarrow E \bullet + E$
5						

Figure 8: Table $\mathcal{T}$ obtained by Earley's algorithm, represented as upper triangular matrix.

The string w is recognized since the final stack symbol $S \rightarrow E \bullet$ is found in $\mathcal{T}_{0,5}$. Observe that $(0, S \rightarrow E \bullet, 5)$ can be derived from $(0, S \rightarrow \bullet E, 0)$ and $(0, E \rightarrow E * E \bullet, 5)$ or from $(0, S \rightarrow \bullet E, 0)$ and $(0, E \rightarrow E + E \bullet, 5)$. This indicates that w is ambiguous.

It can be easily verified that Earley's algorithm adds an item $(j, A \rightarrow \alpha \bullet \beta, i)$ to $\mathcal{T}$ if and only if:

1. $S \Rightarrow^* a_1 \cdots a_j A \gamma$, for some γ , and
2. $\alpha \Rightarrow^* a_{j+1} \cdots a_i$.

In words, the existence of such an item in the table means that there is a derivation from the start symbol S that reaches A , the part of that derivation to the left of that occurrence of A derives the input from position 0 up to position j , and the prefix α of the right-hand side of rule $A \rightarrow \alpha \beta$ derives the input from position j up to position i .

The tabular algorithm of Figure 7 runs in time $\mathcal{O}(|\mathcal{G}|^2 n^3)$ and space $\mathcal{O}(|\mathcal{G}| n^2)$, for a CFG $\mathcal{G}$ and for an input string of length n . Both upper

bounds can be easily derived from the general complexity results discussed in Section 3, taking into account the simplification of items to 3-tuples.

To obtain a formulation of Earley’s algorithm closer to a practical implementation, such as that in Figure 4, read the remarks at the end of Section 3 concerning the agenda and transitions that read the empty string. Alternatively, one may also preprocess certain steps to avoid some of the problems with the agenda during parse time, as discussed by [12], who also showed that the worst-case time complexity of Earley’s algorithm can be improved to $\mathcal{O}(|\mathcal{G}| n^3)$.

5 The Cocke-Kasami-Younger algorithm

Another parsing strategy is (pure) bottom-up parsing, which is also called shift-reduce parsing [33]. It is particularly simple if the CFG $\mathcal{G} = (\Sigma, N, S, R)$ is in Chomsky normal form, which means that each rule is either of the form $A \rightarrow a$, where $a \in \Sigma$, or of the form $A \rightarrow B C$, where $B, C \in N$. The set of stack symbols is the set of nonterminals of the grammar, and the transitions are:

1. $\varepsilon \xrightarrow{a} A$ for each rule $A \rightarrow a$;
2. $B C \xrightarrow{\varepsilon} A$ for each rule $A \rightarrow B C$.

A transition of type 1 consumes the next unread input symbol, and pushes on the stack the nonterminal in the left-hand side of a corresponding rule. A transition of type 2 can be applied if the top-most two stack symbols B and C are such that $B C$ is the right-hand side of a rule, and it replaces B and C by the left-hand side A of that rule. Transitions of types 1 and 2 are called *shift* and *reduce*, respectively; see also Section 6. The final stack symbol is S . We deviate from the other sections in this chapter however by assuming that the PDA starts with an empty stack, or alternatively, that there is some imaginary initial stack symbol that is not in N .

The transitions $B C \xrightarrow{\varepsilon} A$ are of a type that we have seen before, and in a tabular algorithm for the PDA, such transitions can be realized by the inference rule:

1. Let $\mathcal{T} = \emptyset$.
2. For $i = 1, \dots, n$ do:
 - (a) For each rule $A \rightarrow a_i$, add $(i - 1, A, i)$ to $\mathcal{T}$.
 - (b) For $k = i - 2, \dots, 0$ and $j = k + 1, \dots, i - 1$ do:
 - For each rule $A \rightarrow B C$ and all $(k, B, j), (j, C, i) \in \mathcal{T}$, add (k, A, i) to $\mathcal{T}$.
3. Recognize the input if $(0, S, n) \in \mathcal{T}$.

Figure 9: Tabular bottom-up parsing, or the CKY algorithm.

$$\frac{\begin{matrix} (k, B, j) \\ (j, C, i) \end{matrix}}{(k, A, i)} \left\{ \begin{matrix} A \rightarrow BC \end{matrix} \right.$$

Here we use 3-tuples for items, since the first components of the general 4-tuples are redundant, just as in the case of Earley's algorithm above. Transitions of the type $\varepsilon \xrightarrow{a} A$ are new, but they are similar to transitions of the familiar form $B \xrightarrow{a} B A$, where B can be any stack symbol. Because B is irrelevant for deciding whether such a transition can be applied, the expected inference rule

$$\frac{(j, B, i - 1)}{(i - 1, A, i)} \left\{ \begin{matrix} A \rightarrow a_i \\ B \in N \end{matrix} \right.$$

can be simplified to

$$\frac{}{(i - 1, A, i)} \left\{ \begin{matrix} A \rightarrow a_i \end{matrix} \right.$$

A formulation of the tabular bottom-up algorithm closer to a typical implementation is given in Figure 9. This algorithm is also known as the Cocke-Kasami-Younger (CKY) algorithm [42, 2]. Note that no agenda is needed. It can be easily verified that the CKY algorithm adds an item (j, A, i) to $\mathcal{T}$ if and only if $A \Rightarrow^* a_{j+1} \dots a_i$.

As an example, consider the CFG with $\Sigma = \{a, b\}$, $N = \{S, A\}$ and with rules $S \rightarrow SS$, $S \rightarrow AA$, $S \rightarrow b$, $A \rightarrow AS$, $A \rightarrow AA$ and $A \rightarrow a$, and consider

	1	2	3	4
0	A	S, A	S, A	S, A
1		A	A	A
2			S	S
3				S

Figure 10: Table $\mathcal{T}$ obtained by the CKY algorithm.

the input string $w = aabb$. The table $\mathcal{T}$ produced by the CKY algorithm is given in Figure 10, represented as an upper triangular matrix. (Note that the sets $\mathcal{T}_{i,i}$, $0 \leq i \leq n$, on the diagonal of the matrix are always empty and are therefore omitted.) The string w is recognized since the final stack symbol S is found in $\mathcal{T}_{0,4}$.

For a CFG $\mathcal{G} = (\Sigma, N, S, R)$ in Chomsky normal form and an input string of length n , the tabular algorithm of Figure 9 runs in time $\mathcal{O}(|R| n^3)$ and space $\mathcal{O}(|N| n^2)$. Again, these upper bounds can be easily derived from the general complexity results discussed in Section 3, taking into account the simplification of items to 3-tuples. Note that the CKY algorithm runs in time proportional to the size of the grammar, since $|\mathcal{G}| = \mathcal{O}(|R|)$ for CFGs in Chomsky normal form. However, known transformations to Chomsky normal form may increase the size of the grammar by a square function [13].

6 Tabular LR parsing

A more complex parsing strategy is LR parsing [16, 34]. Its main importance is that it results in deterministic PDAs for many practical CFGs for programming languages. For CFGs used in natural language systems however, the resulting PDAs are typically nondeterministic. Although in this case the advantages over simpler parsing strategies have never been convincingly shown, the frequent treatment of nondeterministic LR parsing in recent literature warrants its discussion here.

A distinctive feature of LR parsing is that commitment to a certain rule is postponed until all grammar symbols in the right-hand side of that rule have been found to generate appropriate substrings of the input. In particular, different rules for which this has not yet been accomplished are processed

simultaneously, without spending computational effort on any rule individually. As in the case of Earley's algorithm, we need dotted rules of the form $A \rightarrow \alpha \bullet \beta$, where the dot separates the grammar symbols in the right-hand side that have already been found to derive substrings of the read input from those that are still to be processed. Whereas in the scanner step (3) and in the completer step (4) from Earley's algorithm (Figure 7) each rule is individually processed by letting the dot traverse its right-hand side, in LR parsing this traversal simultaneously affects *sets* of dotted rules. Also the equivalent of the predictor step (2) from Earley's algorithm is now an operation on sets of dotted rules. These operations are pre-compiled into stack symbols and transitions.

Let us fix a CFG $\mathcal{G} = (\Sigma, N, S, R)$. Assume q is a set of dotted rules. We define $\text{closure}(q)$ as the smallest set of dotted rules such that:

1. $q \subseteq \text{closure}(q)$, and
2. if $(A \rightarrow \alpha \bullet B\beta) \in \text{closure}(q)$ and $(B \rightarrow \gamma) \in R$, then $(B \rightarrow \bullet \gamma) \in \text{closure}(q)$.

In words, we extend the set of dotted rules by those that can be obtained by repeatedly applying an operation similar to the predictor step. For a set q of dotted rules and a grammar symbol $X \in \Sigma \cup N$, we define:

$$\text{goto}(q, X) = \text{closure}(\{(A \rightarrow \alpha X \bullet \beta) \mid (A \rightarrow \alpha \bullet X\beta) \in q\})$$

The manner in which the dot traverses through right-hand sides can be related to the scanner step of Earley's algorithm if $X \in \Sigma$ or to the completer step if $X \in N$.

The initial stack symbol q_{init} is defined to be $\text{closure}(\{(S \rightarrow \bullet \alpha) \mid (S \rightarrow \alpha) \in R\})$; cf. the initializer step (1) of Earley's algorithm. Other stack symbols are those non-empty sets of dotted rules that can be derived from q_{init} by means of repeated application of the goto function. More precisely, Q is the smallest set such that:

1. $q_{\text{init}} \in Q$, and
2. if $q \in Q$ and $\text{goto}(q, X) = q' \neq \emptyset$ for some X , then $q' \in Q$.

For technical reasons, we also need to add a special stack symbol q_{final} to Q that becomes the final stack symbol. The transitions are:

1. $q_1 \xrightarrow{a} q_1 q_2$ for all $q_1, q_2 \in Q$ and each $a \in \Sigma$ such that $\text{goto}(q_1, a) = q_2$;
2. $q_0 q_1 \cdots q_m \xrightarrow{\varepsilon} q_0 q'$ for all $q_0, \dots, q_m, q' \in Q$ and each $(A \rightarrow \alpha \bullet) \in q_m$ such that $|\alpha| = m$ and $q' = \text{goto}(q_0, A)$;
3. $q_0 q_1 \cdots q_m \xrightarrow{\varepsilon} q_{\text{final}}$ for all $q_0, \dots, q_m \in Q$ and each $(S \rightarrow \alpha \bullet) \in q_m$ such that $|\alpha| = m$ and $q_0 = q_{\text{init}}$.

The first type of transition is called *shift*. It can be seen as the pre-compilation of the scanner step followed by repeated application of the predictor step. Note that only one transition is applied for each input symbol that is read, independent of the number of dotted rules in the sets q_1 and q_2 . The second type of transition is called *reduction*. It can be applied when the symbol on top of the stack contains a dotted rule with the dot at the end of the right-hand side. First, as many symbols are popped from the stack as that right-hand side is long, and then a symbol $q' = \text{goto}(q_0, A)$ is pushed on the stack. This is related to the completer step from Earley's algorithm. The third type of transition is very similar to the second. It is only applied once, when the start symbol has been found to generate (a prefix of) the input.

For tabular LR parsing, we apply the same framework as in the previous sections, to obtain Figure 11. A slight difficulty is caused by the new types of transition $q_0 \cdots q_m \xrightarrow{\varepsilon} q_0 q'$ and $q_0 \cdots q_m \xrightarrow{\varepsilon} q_{\text{final}}$, but these can be handled by a straightforward generalization of inference rules from Figures 5 and 6. Note that we need 4-tuple items here rather than 3-tuple items as in the previous two sections. Tabular LR parsing is also known as *generalized LR parsing* [37, 38]. In the literature on generalized LR parsing, but only there, the table $\mathcal{T}$ of items is often called a *graph-structured stack*.

As an example, consider the grammar with the rules $S \rightarrow S + S$ and $S \rightarrow a$. Apart from q_{final} , the stack symbols of the PDA are represented in Figure 12 as rectangles enclosing sets of dotted rules. There is an arc from stack symbol q to stack symbol q' labelled by X to denote that $\text{goto}(q, X) = q'$. For the input $a + a + a$, the table $\mathcal{T}$ is given by Figure 13. Note that q_3 at position 4 has two outgoing arcs, since it can arise by a shift with $+$ from q_4 or from q_2 . Also note that $(\perp, 0, q_{\text{final}}, 5)$ is found twice, once from $(\perp, 0, q_{\text{init}}, 0)$, $(q_{\text{init}}, 0, q_2, 1)$, $(q_2, 1, q_3, 2)$, $(q_3, 2, q_4, 5)$, and once from $(\perp, 0, q_{\text{init}}, 0)$, $(q_{\text{init}}, 0, q_2, 3)$, $(q_2, 3, q_3, 4)$, $(q_3, 4, q_4, 5)$, in both cases by means of $(S \rightarrow S + S \bullet) \in q_4$, with $|S + S| = 3$. This indicates that the input $a + a + a$ is ambiguous.

$$\overline{(\perp, 0, q_{init}, 0)} \quad (5)$$

$$\frac{(q', j, q_1, i-1)}{(q_1, i-1, q_2, i)} \left\{ \begin{array}{l} goto(q_1, a_i) = q_2 \end{array} \right. \quad (6)$$

$$\frac{\begin{array}{c} (q_0, j_0, q_1, j_1) \\ (q_1, j_1, q_2, j_2) \\ \vdots \\ (q_{m-1}, j_{m-1}, q_m, j_m) \end{array}}{(q_0, j_0, q', j_m)} \left\{ \begin{array}{l} (A \rightarrow \alpha \bullet) \in q_m \\ |\alpha| = m \\ q' = goto(q_0, A) \end{array} \right. \quad (7)$$

$$\frac{\begin{array}{c} (\perp, 0, q_0, j_0) \\ (q_0, j_0, q_1, j_1) \\ (q_1, j_1, q_2, j_2) \\ \vdots \\ (q_{m-1}, j_{m-1}, q_m, j_m) \end{array}}{(\perp, 0, q_{final}, j_m)} \left\{ \begin{array}{l} (S \rightarrow \alpha \bullet) \in q_m \\ |\alpha| = m \\ q_0 = q_{init} \end{array} \right. \quad (8)$$

Figure 11: Tabular LR parsing, or generalized LR parsing.

If the grammar at hand does not contain rules of the form $A \rightarrow \varepsilon$, then the tabular algorithm from Figure 11 can be reformulated in a way very similar to the algorithm from Figure 4. If there are rules of the form $A \rightarrow \varepsilon$ however, the handling of the agenda is complicated, due to problems similar to those we discussed at the end of Section 3. This issue is investigated by [28, 25].

We now analyze the time and space complexity of tabular LR parsing. Let us fix a CFG $\mathcal{G} = (\Sigma, N, S, R)$. Let p be the length of the longest right-hand side of a rule in R and let n be the length of the input string. Once again, we assume that $\mathcal{T}$ is implemented as a square array of size $n+1$. Consider the reduction step (7) in Figure 11. Each application of this step is uniquely identified by $m+1 \leq p+1$ input positions and $|Q||R|$




Figure 12: The set of stack symbols, excluding q_{final} , and the *goto* function.




Figure 13: Table $\mathcal{T}$ obtained by tabular LR parsing.

combinations of stack symbols. The expression $|Q| |R|$ is due to the fact that, once a stack symbol q_0 and a rule $A \rightarrow X_1 X_2 \cdots X_m$ have been selected such that $(A \rightarrow \bullet X_1 X_2 \cdots X_m) \in q_0$, then stack symbols q_i , $1 \leq i \leq m$, and q' are uniquely determined by $q_1 = \text{goto}(q_0, X_1)$, $q_2 = \text{goto}(q_1, X_2)$, $\dots$, $q_m = \text{goto}(q_{m-1}, X_m)$ and $q' = \text{goto}(q_0, A)$. (As can be easily verified, a derivable stack of which the top-most symbol q_m contains $(A \rightarrow X_1 X_2 \cdots X_m \bullet)$ must necessarily have top-most symbols $q_0 q_1 \cdots q_m$ with the above constraints.) Since a single application of this step can easily be carried out in time $\mathcal{O}(p)$, we conclude the total amount of time required by all applications of the step is $\mathcal{O}(|Q| |R| p n^{p+1})$. This is also the worst-case time complexity of the

algorithm, since the running time is dominated by the reduction step (7). From the general complexity results discussed in Section 3 it follows that the worst-case space complexity is $\mathcal{O}(|Q|^2 n^2)$.

We observe that while the above time bound is polynomial in the length of the input string, it can be much worse than the corresponding bounds for Earley's algorithm or for the CKY algorithm, since p is not bounded. A solution to this problem has been discussed by [15, 26] and consists in splitting each reduction into $\mathcal{O}(p)$ transitions of the form $q'q'' \xrightarrow{\varepsilon} q$. In this way, the maximum length of transitions becomes independent of the grammar. This results in tabular implementations of LR parsing with cubic time complexity in the length of the input. We furthermore observe that the term $|Q|$ in the above bounds depends on the specific structure of $\mathcal{G}$, and may grow exponentially with $|\mathcal{G}|$ [34, Proposition 6.46].

7 Parse trees

As stated in Section 1, recognition is the process of determining whether an input string is in the language described by a grammar or automaton, and parsing is the process of determining the parse trees of an input string according to a grammar. Although the algorithms we have discussed up to now are recognition algorithms, they can be easily extended to become parsing algorithms, as we show in this section. In what follows we assume a fixed CFG $\mathcal{G} = (\Sigma, N, S, R)$ and an input string $w = a_1 \cdots a_n \in \Sigma^*$.

Since the number of parse trees can be exponential in the length of the input string, and even infinite when $\mathcal{G}$ is cyclic, one first needs to find a way to compactly represent the set of all parse trees. This is usually done through a CFG $\mathcal{G}_w$, called *parse forest*, defined as follows. The alphabet of $\mathcal{G}_w$ is the same as that of $\mathcal{G}$, and the nonterminals of $\mathcal{G}_w$ have the form (j, A, i) , where $A \in N$ and $0 \leq j \leq i \leq n$. The start symbol of $\mathcal{G}_w$ is $(0, S, n)$. The rules of $\mathcal{G}_w$ include at least those of the form $(i_0, A, i_m) \rightarrow (i_0, X_1, i_1) \cdots (i_{m-1}, X_m, i_m)$, where (i) $(A \rightarrow X_1 \cdots X_m) \in R$, (ii) $S \Rightarrow^* a_1 \cdots a_{i_0} A a_{i_m+1} \cdots a_n$, and (iii) $X_j \Rightarrow^* a_{i_{j-1}+1} \cdots a_{i_j}$ for $1 \leq j \leq m$, and those of the form $(i-1, a_i, i) \rightarrow a_i$. However, $\mathcal{G}_w$ may also contain rules $(i_0, A, i_m) \rightarrow (i_0, X_1, i_1) \cdots (i_{m-1}, X_m, i_m)$ that violate constraints (ii) or (iii) above. Such rules cannot be part of any derivation of a terminal string from $(0, S, n)$ and can be eliminated by a process that is called *reduction*.

Reduction can be carried out in linear time in the size of $\mathcal{G}_w$ [33].

It is not difficult to show that the parse forest $\mathcal{G}_w$ generates a finite language, which is either $\{w\}$ if w is in the language generated by $\mathcal{G}$, or $\emptyset$ otherwise. Furthermore, there is a one-to-one correspondence between parse trees according to $\mathcal{G}_w$ and parse trees of w according to $\mathcal{G}$, with corresponding parse trees being isomorphic.

To give a concrete example, let us consider the CKY algorithm presented in Section 5. In order to extend this recognition algorithm to a parsing algorithm, we may construct the parse forest $\mathcal{G}_w$ with rules of the form $(j, A, i) \rightarrow (j, B, k) (k, C, i)$, where $(A \rightarrow B C) \in R$ and $(j, B, k), (k, C, i) \in \mathcal{T}$, rules of the form $(i-1, A, i) \rightarrow (i-1, a_i, i)$, where $(A \rightarrow a_i) \in R$, and rules of the form $(i-1, a_i, i) \rightarrow a_i$. Such rules can be constructed during the computation of the table $\mathcal{T}$. In order to perform reduction on $\mathcal{G}_w$, one may visit the nonterminals of $\mathcal{G}_w$ starting from $(0, S, n)$, following the rules in a top-down fashion, eliminating the nonterminals and the associated rules that are never reached. From the resulting parse forest $\mathcal{G}_w$, individual parse trees can be extracted in time proportional to the size of the parse tree itself, which in the case of CFGs in Chomsky normal form is $\mathcal{O}(n)$. One may also extract parse trees directly from table $\mathcal{T}$, but the time complexity then becomes $\mathcal{O}(|\mathcal{G}| n^2)$ [2, 11].

Consider the table $\mathcal{T}$ from Figure 10, which was produced by the CKY algorithm with $w = aabb$ and $\mathcal{G} = (\Sigma, N, S, R)$, where $\Sigma = \{a, b\}$, $N = \{S, A\}$ and $R = \{S \rightarrow SS, S \rightarrow AA, S \rightarrow b, A \rightarrow AS, A \rightarrow AA, A \rightarrow a\}$. The method presented above constructs the parse forest $\mathcal{G}_w = (\Sigma, N_w, (0, S, 4), R_w)$, where $N_w \subseteq \{(j, B, i) \mid B \in N, 0 \leq j < i \leq 4\}$ and R_w contains the rules in Figure 14. Rules that are eliminated by reduction are marked by $\dagger$.

If $\mathcal{G}$ is in Chomsky normal form, then we have $|\mathcal{G}_w| = \mathcal{O}(|\mathcal{G}| n^3)$. For general CFGs, however, we have $|\mathcal{G}_w| = \mathcal{O}(|\mathcal{G}| n^{p+1})$, where p is the length of the longest right-hand side of a rule in $\mathcal{G}$. In practical parsing applications this higher space complexity is usually avoided by applying the following method, which is based on [17, 6]. In place of computing $\mathcal{G}_w$, one constructs an alternative CFG containing rules of the form $t \rightarrow t_1 \cdots t_m$, where $t, t_1, \dots, t_m \in \mathcal{T}$ such that item t was derived from items $t_1, \dots, t_m$ via an inference rule with m antecedents. Parse trees according to this new CFG can be extracted as usual. From these trees, the desired parse trees for w according to $\mathcal{G}$

$(0, a, 1) \rightarrow a$	$(0, S, 3) \rightarrow (0, A, 1) (1, A, 3)$
$(1, a, 2) \rightarrow a$	$(0, S, 3) \rightarrow (0, S, 2) (2, S, 3)$
$(2, b, 3) \rightarrow b$	$(0, A, 3) \rightarrow (0, A, 1) (1, A, 3) \dagger$
$(3, b, 4) \rightarrow b$	$(0, A, 3) \rightarrow (0, A, 2) (2, S, 3) \dagger$
$(0, A, 1) \rightarrow (0, a, 1)$	$(1, A, 4) \rightarrow (1, A, 2) (2, S, 4)$
$(1, A, 2) \rightarrow (1, a, 2)$	$(1, A, 4) \rightarrow (1, A, 3) (3, S, 4)$
$(2, S, 3) \rightarrow (2, b, 3)$	$(0, S, 4) \rightarrow (0, A, 1) (1, A, 4)$
$(3, S, 4) \rightarrow (3, b, 4)$	$(0, S, 4) \rightarrow (0, S, 2) (2, S, 4)$
$(0, S, 2) \rightarrow (0, A, 1) (1, A, 2)$	$(0, S, 4) \rightarrow (0, S, 3) (3, S, 4)$
$(0, A, 2) \rightarrow (0, A, 1) (1, A, 2) \dagger$	$(0, A, 4) \rightarrow (0, A, 1) (1, A, 4) \dagger$
$(1, A, 3) \rightarrow (1, A, 2) (2, S, 3)$	$(0, A, 4) \rightarrow (0, A, 2) (2, S, 4) \dagger$
$(2, S, 4) \rightarrow (2, S, 3) (3, S, 4)$	$(0, A, 4) \rightarrow (0, A, 3) (3, S, 4) \dagger$

Figure 14: Parse forest associated with table $\mathcal{T}$ from Figure 10.

can be easily obtained by elementary tree editing operations such as node relabelling and node erasing. The precise editing algorithm that should be applied depends on the deduction system underlying the adopted recognition algorithm.

If the adopted recognition algorithm has inference rules with no more than $m = 2$ antecedents, then the space complexity of the parsing method discussed above, expressed as a function of the length n of the input string, is $\mathcal{O}(n^3)$. Note that $m = 2$ in the case of Earley's algorithm, and this also holds in practical implementations of tabular LR parsing, as discussed at the end of Section 6. The space complexity in the size of $\mathcal{G}$ may be larger than $\mathcal{O}(|\mathcal{G}|)$, however; it is $\mathcal{O}(|\mathcal{G}|^2)$ in the case of Earley's algorithm and even exponential in the case of tabular LR parsing.

The parse forest representation is originally due to [5], with states of a finite automaton in place of positions in an input string. Parse forests have also been discussed by [7, 30, 37, 22]. Similar ideas were proposed for tree-adjointing grammars by [40, 18] (see [14] and references therein for the definition of tree-adjointing grammars).

8 Further references

In this chapter we have restricted ourselves to tabulation for context-free parsing, on the basis of PDAs. A similar kind of tabulation was also developed for tree-adjoining grammars on the basis of an extended type of PDA [3]. Tabulation for an even more general type of PDA was discussed by [41].

A further restriction we have made is that the input to the parser must be a string. Context-free parsing can however be generalized to input consisting of a finite automaton. Finite automata without cycles used in speech recognition systems are also referred to as word graphs or word lattices [4]. The parsing methods developed in this chapter can be easily adapted to parsing of finite automata, by manipulating states of an input automaton in place of positions in an input string. This technique can be traced back to [5], which we mentioned before in Section 7.

PDAs are usually considered to read input from left to right, and the forms of tabulation that we discussed follow that directionality.² For types of tabular parsing that are not strictly in one direction, such as head-driven parsing [32] and island-driven parsing [29], it is less appealing to take PDAs as starting point.

Earley's algorithm and the CKY algorithm run in cubic time in the length of the input string. An asymptotically faster method for context-free parsing has been developed by [39], using a reduction from context-free recognition to Boolean matrix multiplication. An inverse reduction from Boolean matrix multiplication to context-free recognition has been presented by [20], providing evidence that asymptotically faster methods for context-free recognition might not be of practical interest.

The extension of tabular parsing with weights or probabilities has been considered by [23] for Earley's algorithm, by [35] for the CKY algorithm, and by [19] for tabular LR parsing. Deduction systems for parsing extended with weights are discussed by [10].

²There are alternative forms of tabulation that do not adopt the left-to-right mode of processing from the PDA [1, 24].

References

- [1] A.V. Aho, J.E. Hopcroft, and J.D. Ullman. Time and tape complexity of pushdown automaton languages. *Information and Control*, 13:186–206, 1968.
- [2] A.V. Aho and J.D. Ullman. *Parsing*, volume 1 of *The Theory of Parsing, Translation and Compiling*. Prentice-Hall, 1972.
- [3] M. A. Alonso Pardo, M.-J. Nederhof, and E. Villemonte de la Clergerie. Tabulation of automata for tree-adjointing languages. *Grammars*, 3:89–110, 2000.
- [4] H. Aust, M. Oerder, F. Seide, and V. Steinbiss. The Philips automatic train timetable information system. *Speech Communication*, 17:249–262, 1995.
- [5] Y. Bar-Hillel, M. Perles, and E. Shamir. On formal properties of simple phrase structure grammars. In Y. Bar-Hillel, editor, *Language and Information: Selected Essays on their Theory and Application*, chapter 9, pages 116–150. Addison-Wesley, 1964.
- [6] S. Billot and B. Lang. The structure of shared forests in ambiguous parsing. In *27th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference*, pages 143–151, Vancouver, British Columbia, Canada, June 1989.
- [7] J. Cocke and J.T. Schwartz. *Programming Languages and Their Compilers — Preliminary Notes*, pages 184–206. Courant Institute of Mathematical Sciences, New York University, second revised version, April 1970.
- [8] S.A. Cook. Path systems and language recognition. In *ACM Symposium on Theory of Computing*, pages 70–72, 1970.
- [9] J. Earley. An efficient context-free parsing algorithm. *Communications of the ACM*, 13(2):94–102, February 1970.
- [10] J. Goodman. Semiring parsing. *Computational Linguistics*, 25(4):573–605, 1999.

- [11] S.L. Graham and M.A. Harrison. Parsing of general context free languages. In *Advances in Computers*, volume 14, pages 77–185. Academic Press, New York, NY, 1976.
- [12] S.L. Graham, M.A. Harrison, and W.L. Ruzzo. An improved context-free recognizer. *ACM Transactions on Programming Languages and Systems*, 2(3):415–462, July 1980.
- [13] M.A. Harrison. *Introduction to Formal Language Theory*. Addison-Wesley, 1978.
- [14] A.K. Joshi and Y. Schabes. Tree-adjoining grammars. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages. Vol 3: Beyond Words*, chapter 2, pages 69–123. Springer-Verlag, Berlin/Heidelberg/New York, 1997.
- [15] J.R. Kipps. GLR parsing in time $\mathcal{O}(n^3)$. In M. Tomita, editor, *Generalized LR Parsing*, chapter 4, pages 43–59. Kluwer Academic Publishers, 1991.
- [16] D.E. Knuth. On the translation of languages from left to right. *Information and Control*, 8:607–639, 1965.
- [17] B. Lang. Deterministic techniques for efficient non-deterministic parsers. In *Automata, Languages and Programming, 2nd Colloquium*, volume 14 of *Lecture Notes in Computer Science*, pages 255–269, Saarbrücken, 1974. Springer-Verlag.
- [18] B. Lang. Recognition can be harder than parsing. *Computational Intelligence*, 10(4):486–494, 1994.
- [19] A. Lavie and M. Tomita. GLR* – an efficient noise-skipping parsing algorithm for context free grammars. In *Third International Workshop on Parsing Technologies*, pages 123–134, Tilburg (The Netherlands) and Durbuy (Belgium), August 1993.
- [20] L. Lee. Fast context-free grammar parsing requires fast boolean matrix multiplication. *Journal of the ACM*, 49(1):1–15, 2001.
- [21] R. Leermakers. *The Functional Treatment of Parsing*. Kluwer Academic Publishers, 1993.

- [22] H. Leiss. On Kilbury’s modification of Earley’s algorithm. *ACM Transactions on Programming Languages and Systems*, 12(4):610–640, October 1990.
- [23] G. Lyon. Syntax-directed least-errors analysis for context-free languages: A practical approach. *Communications of the ACM*, 17(1):3–14, January 1974.
- [24] M.-J. Nederhof. Reversible pushdown automata and bidirectional parsing. In J. Dassow, G. Rozenberg, and A. Salomaa, editors, *Developments in Language Theory II*, pages 472–481. World Scientific, Singapore, 1996.
- [25] M.-J. Nederhof and J.J. Sarbo. Increasing the applicability of LR parsing. In H. Bunt and M. Tomita, editors, *Recent Advances in Parsing Technology*, chapter 3, pages 35–57. Kluwer Academic Publishers, 1996.
- [26] M.-J. Nederhof and G. Satta. Efficient tabular LR parsing. In *34th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference*, pages 239–246, Santa Cruz, California, USA, June 1996.
- [27] M.J. Nederhof. *Linguistic Parsing and Program Transformations*. PhD thesis, University of Nijmegen, 1994.
- [28] R. Nozohoor-Farshi. GLR parsing for ε -grammars. In M. Tomita, editor, *Generalized LR Parsing*, chapter 5, pages 61–75. Kluwer Academic Publishers, 1991.
- [29] G. Satta and O. Stock. Bidirectional context-free grammar parsing for natural language processing. *Artificial Intelligence*, 69:123–164, 1994.
- [30] B.A. Sheil. Observations on context-free parsing. *Statistical Methods in Linguistics*, pages 71–109, 1976.
- [31] S.M. Shieber, Y. Schabes, and F.C.N. Pereira. Principles and implementation of deductive parsing. *Journal of Logic Programming*, 24:3–36, 1995.
- [32] K. Sikkel. *Parsing Schemata*. Springer-Verlag, 1997.

- [33] S. Sippu and E. Soisalon-Soininen. *Parsing Theory, Vol. I: Languages and Parsing*, volume 15 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, 1988.
- [34] S. Sippu and E. Soisalon-Soininen. *Parsing Theory, Vol. II: LR(k) and LL(k) Parsing*, volume 20 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, 1990.
- [35] R. Teitelbaum. Context-free error analysis by evaluation of algebraic power series. In *Conference Record of the Fifth Annual ACM Symposium on Theory of Computing*, pages 196–199, 1973.
- [36] H. Thompson and G. Ritchie. Implementing natural language parsers. In T. O’Shea and M. Eisenstadt, editors, *Artificial Intelligence: Tools, Techniques, and Applications*, chapter 9, pages 245–300. Harper & Row, New York, 1984.
- [37] M. Tomita. *Efficient Parsing for Natural Language*. Kluwer Academic Publishers, 1986.
- [38] M. Tomita. An efficient augmented-context-free parsing algorithm. *Computational Linguistics*, 13:31–46, 1987.
- [39] L.G. Valiant. General context-free recognition in less than cubic time. *Journal of Computer and System Sciences*, 10:308–315, 1975.
- [40] K. Vijay-Shanker and D.J. Weir. The use of shared forests in tree adjoining grammar parsing. In *Sixth Conference of the European Chapter of the Association for Computational Linguistics, Proceedings of the Conference*, pages 384–393, Utrecht, The Netherlands, April 1993.
- [41] E. Villemonte de la Clergerie and F. Barthélemy. Information flow in tabular interpretations for generalized push-down automata. *Theoretical Computer Science*, 199:167–198, 1998.
- [42] D.H. Younger. Recognition and parsing of context-free languages in time n^3 . *Information and Control*, 10:189–208, 1967.

Tabular Parsing

Giorgio Satta
University of Padua
<http://www.dei.unipd.it/~satta/>

Introduction

Parsing : retrieving all derivations/trees assigned by input grammar G to input string w

Recognition : detecting whether input grammar G generates string w

Recognition and Parsing algorithms are closely related

Introduction

Tabular approaches implement so called **dynamic programming** techniques : input instance is broken into all possible subinstances of the problem, these are solved and the solutions are stored and recombined to obtain solutions of the original instance

Tabular techniques allow “representation sharing”, in contrast with approaches that explore one solution at a time and use backtracking

Introduction

Most popular tabular parsing algorithms for context-free grammars (CFGs) :

- CKY algorithm (Younger, 1967)
- Earley algorithm (Earley, 1970)
- GHR algorithm (Graham, Harrison & Ruzzo, 1980)
- Chart parsing (Kay, 1980)
- Left-corner parsing (Rosenkrantz & Lewis, 1970)
- Generalized LR (Tomita, 1986)

Introduction

Questions often asked by students :

- How do Shift-Reduce, Earley, Left-Corner and LR parsing strategies compare to one another ?
- What is the relation between the chart and the graph-structured stack ? Can we do LR parsing with a chart?
- How can we design a new parsing strategy and implement it using tabular techniques ?

Relations among tabular parsing algorithms are not well documented in the existing literature

Introduction


We present a framework for specification of CFG tabular parsers that can be used to (Lang, 1974) :

- investigate deep similarities/differences between existing algorithms
- design new algorithms
- prove formal properties

The framework is based on the notions of

- push-down automata
- tabular simulation of nondeterminism

Introduction




ESSLLI 2004

Tabular Parsing

7

Introduction



ESSLLI 2004

Tabular Parsing

8

Formal Background

Binary relations :

- $R \subseteq \Delta \times \Delta =_{\text{def}} \{ (d_1, d_2) \mid d_1, d_2 \in \Delta \}$
- $R \cdot R' =_{\text{def}} \{ (d_1, d_2) \mid (d_1, d') \in R, (d', d_2) \in R' \}$
- $R^0 =_{\text{def}} \{ (d, d) \mid d \in \Delta \}$
- $R^{i+1} =_{\text{def}} R \cdot R^i, i \geq 0$
- $R^* =_{\text{def}} \bigcup_{i \geq 0} R^i$
- $R^+ =_{\text{def}} \bigcup_{i \geq 1} R^i$

Formal Background

Alphabets and strings :

- for a string x , $|x|$ denotes the length of x
- ε denotes the empty string, $|\varepsilon| = 0$
- Σ^* denotes the set of all strings over alphabet Σ
- $x \cdot y$ (also written xy) is the concatenation of strings x and y

Indices :

- j is an input position, or **index**, for string x if $0 \leq j \leq |x|$

Formal Background

Languages :

- $L \subseteq \Sigma^*$
- $L \cdot L' =_{\text{def}} \{ x \cdot y \mid x \in L, y \in L' \}$
- $L^0 =_{\text{def}} \{ \varepsilon \}$
- $L^{i+1} =_{\text{def}} L \cdot L^i, i \geq 0$
- $R^* =_{\text{def}} \bigcup_{i \geq 0} L^i$
- $R^+ =_{\text{def}} \bigcup_{i \geq 1} L^i$

Formal Background

Complexity :

- $f(n) = O(g(n))$ iff there are n_0 and c such that, for every $n \geq n_0$, we have $0 \leq f(n) \leq c \cdot g(n)$
- $f(n) = \Theta(g(n))$ iff there are n_0 , c and c' such that, for every $n \geq n_0$, we have $0 \leq c \cdot g(n) \leq f(n) \leq c' \cdot g(n)$

Context-Free Grammars

Context-Free Grammar (CFG) :

$$G =_{\text{def}} (\Sigma, N, P, S)$$

- Σ terminal symbols
- N nonterminal symbols
- $S \in N$ start symbol
- P set of productions having the form $A \rightarrow \alpha$, with $A \in N$ and $\alpha \in (N \cup \Sigma)^*$
- CFG is the class of all context-free grammars

Context-Free Grammars

Conventions :

- $V = N \cup \Sigma$
- $A, B, C, \dots \in N$
- $a, b, c, \dots \in \Sigma$
- $X, Y, \dots \in \Sigma \cup \{\varepsilon\}$
- $\alpha, \beta, \gamma, \dots \in V^*$
- $u, v, w, x, \dots \in \Sigma^*$

Context-Free Grammars

Derivations :

- $\gamma A \delta \Rightarrow \gamma \alpha \delta$, if $p = (A \rightarrow \alpha) \in P$; also written as $\gamma A \delta \Rightarrow^p \gamma \alpha \delta$
- $x A \delta \Rightarrow_{\text{lm}} x \alpha \delta$, if $p = (A \rightarrow \alpha) \in P$
- we also use $\Rightarrow^*$ and $\Rightarrow^+$
- $\gamma \Rightarrow^\pi \delta$, $\pi = p_1 p_2 \dots p_q \in P^*$, if we have $\gamma = \gamma_0 \Rightarrow^{p^1} \gamma_1 \Rightarrow^{p^2} \gamma_2 \dots \gamma_{q-1} \Rightarrow^{p^q} \gamma_q = \delta$
- $S \Rightarrow^\pi w$ is called a **sentential** derivation

Context-Free Grammars

Derivations :

- $L(G) =_{\text{def}} \{ w \mid S \Rightarrow^* w \}$
- CFL is the class of all languages generated by CFGs
- G is **reduced** if, for every $A \in N$, $S \Rightarrow^* \gamma A \delta \Rightarrow w$
- G is **cyclic** if G is reduced and, for some $A \in N$, $A \Rightarrow^+ A$

Context-Free Grammars

Parse trees :

- sentential derivations can be described by means of labeled, ordered trees called **parse trees**
- parse trees do not represent the order of rule application, but represent the nonterminal that is rewritten by each rule

ESSLI 2004


Tabular Parsing

17

Context-Free Grammars

Example :

- $G = (\{S\}, \{a, b\}, P, S)$
- $P = \{S \rightarrow a S b, S \rightarrow a b\}$



ESSLI 2004

Tabular Parsing

18

Context-Free Grammars

Parse trees :

- $T(G) =_{\text{def}}$ set of all parse trees for G
- $T(G, w) =_{\text{def}}$ set of all parse trees for G with yield w

Context-Free Grammars

Ambiguity :

- G is **unambiguous** if, for every $w \in L(G)$ we have $|T(G, w)| = 1$
- G is **ambiguous** if it is not unambiguous
- $L \in \text{CFL}$ is **ambiguous** if there is no CFG G such that $L(G) = L$ and G is unambiguous

Context-Free Grammars

Ambiguity :

- $d(n) =_{\text{def}} \max_{|w| = n} |T(G, w)|$
- G is **massively ambiguous** if there is no polynomial $p(n)$ such that $d(n) = O(p(n))$
- G is **infinitely ambiguous** if, for some w , $T(G, w)$ is an infinite set

Context-Free Grammars

Example :

- $G = (\{S\}, \{a\}, P, S)$
- $P = \{S \rightarrow SS, S \rightarrow a\}$
- G is massively ambiguous, generating all possible binary bracketings for a string
- $d(n)$ is the series of Catalan numbers

Context-Free Grammars

Example :

- $G = (\Sigma, N, P, S)$
- $N = \{ S, VP, NP, PP, V, N, P, Det \}$
- $\Sigma = \{ \text{eat, I, chocolate, fork, strawberries, with, a} \}$
- $P = \{ \begin{array}{l} NP \rightarrow NP PP \mid Det N \mid N, \\ VP \rightarrow VP PP \mid V NP, \quad PP \rightarrow P NP, \\ N \rightarrow \text{chocolate} \mid I \mid \text{fork} \mid \text{strawberries}, \\ V \rightarrow \text{eat}, \quad Det \rightarrow a, \quad P \rightarrow \text{with} \end{array} \}$

Context-Free Grammars

Example :

- I eat $[_{NP} [_{NP} \text{chocolate}] [_{PP} \text{with strawberries}]]$
- I $[_{VP} [_{VP} \text{eat} [_{NP} \text{chocolate}]] [_{PP} \text{with a fork}]]$
- *I $[_{VP} [_{VP} \text{eat} [_{NP} \text{chocolate}]] [_{PP} \text{with strawberries}]]$
- *I eat $[_{NP} [_{NP} \text{chocolate}] [_{PP} \text{with a fork}]]$

Context-Free Grammars

Complexity :

- $|G| =_{\text{def}} \sum_{(A \rightarrow \alpha) \in P} |A \cdot \alpha|$
- the complexity of a CFG is a reasonable measure of the size we need to represent the grammar within a machine

Push-down automata

Push-Down Automaton (PDA) :

$$M =_{\text{def}} (\Sigma, Q, q_{in}, q_{fin}, \Delta)$$

- Σ input symbols
- Q stack symbols
- $q_{in}, q_{fin} \in Q$ initial and final stack symbols, resp.
- Δ set of transitions having the form $\sigma_1 \Rightarrow_X \sigma_2$, with $\sigma_1, \sigma_2 \in Q$ and $X \in \Sigma \cup \{\varepsilon\}$
- note that there is no internal state in M , which can be encoded into Q without blowing up the representation

Push-down automata

Configuration :

- fix an input string $w = a_1 a_2 \dots a_n$, $a_i \in \Sigma$
- a configuration of M on w is a pair (σ, i) , where
 - $\sigma \in Q^*$ is a stack, with top-most symbol at the right
 - $0 \leq i \leq n$ is an input position, indicating the number of symbols already read so far

Push-down automata

Move :

- we write $(\sigma, j) \Rightarrow (\sigma', i)$ iff
 - $(\sigma_1 \Rightarrow_X \sigma_2) \in \Delta$
 - $\sigma = \sigma'' \sigma_1$, $\sigma' = \sigma'' \sigma_2$, $\sigma'' \in Q^*$
 - $X = a_{j+1}$ and $j = i+1$, or else $X = \varepsilon$ and $j = i$
- $L(M) =_{\text{def}} \{ w \mid (q_{in}, 0) \Rightarrow (q_{fin}, |w|) \}$

Push-down automata

Example :

- $\Sigma = \{ a, b, c, d \}$
- $Q = \{ q_0, \dots, q_9 \}$, q_0 initial, q_9 final
- Δ :

$q_0 \Rightarrow_a q_0 q_1$	$q_4 q_5 \Rightarrow_\epsilon q_6$
$q_0 q_1 \Rightarrow_b q_0 q_2$	$q_2 q_6 \Rightarrow_\epsilon q_7$
$q_0 q_1 \Rightarrow_b q_0 q_3$	$q_0 q_7 \Rightarrow_\epsilon q_9$
$q_2 \Rightarrow_c q_2 q_4$	$q_3 q_6 \Rightarrow_\epsilon q_8$
$q_3 \Rightarrow_c q_3 q_4$	$q_0 q_8 \Rightarrow_\epsilon q_9$
$q_4 \Rightarrow_d q_4 q_5$	


ESSLLI 2004

Tabular Parsing

29

Push-down automata

Example : $w = a b c d$




ESSLLI 2004

Tabular Parsing

30

Push-down automata

Example (cont'd) :



Push-down automata

Nondeterminism :

- M is **deterministic** if for each input string and each configuration there is at most one transition that can be applied
- M is **nondeterministic** if it is not deterministic

Push-down automata

Complexity :

- $|M| =_{\text{def}} \sum_{(\sigma_1 \Rightarrow_X \sigma_2) \in \Delta} |\sigma_1 \cdot X \cdot \sigma_2|$
- the complexity of a PDA is a reasonable measure of the size we need to represent the automaton within a machine

Tabulation


We want to simulate a nondeterministic PDA :

- there can be exponentially many computations on given input string
- some computations might not even terminate

The problem can be solved using dynamic programming to obtain [computation sharing](#)

Tabulation

For given j, i, q and q' , consider the following computation chunks :



The overall effect of the above computations is to “push” q' on top of q when processing the input from j to i

ESSLLI 2004

Tabular Parsing


35

Tabulation

We group all these chunks into an equivalence class

We assume a graph, called **chart**, whose vertices represent indices of the input string

The equivalence class is uniquely represented as an arc spanning indices j and i , with label (q, q')



ESSLLI 2004

Tabular Parsing

36

Tabulation

For technical reasons, we use a special symbol $\perp$, which we assume to occur below the bottom of the stack; this is needed to represent stacks with a single symbol

We always start with the following arc in our chart, denoting the initial stack symbol :




ESSLLI 2004

Tabular Parsing

37

Tabulation

Example (cont'd) :




ESSLLI 2004

Tabular Parsing

38


Tabulation

Example (cont'd) :



Tabulation

Example (cont'd) :



Tabulation

Assume a path is constructed from vertex 0 to vertex i with arc labels

$$(\perp, q_1), (q_1, q_2), \dots, (q_{m-1}, q_m)$$


Then there exists at least one computation of the PDA having the form

$$(q_{in}, 0) \Rightarrow^* (q_1 q_2 \dots q_{m-1} q_m, i)$$

Tabulation

Path : $(\perp, q_1), (q_1, q_2), \dots, (q_{m-1}, q_m)$

Computation(s) :



Tabulation

Tabulation algorithm :

- $w = a_1 a_2 \dots a_n, n \geq 0$
- $I = \{ (q', j, q, i) \mid 0 \leq j \leq i \leq n, q', q \in Q \}$
= set of **items**
- T = **table** of computed items
- N = **agenda** of items to be combined

Tabulation

The following statements are equivalent :

- item (q', j, q, i) is added to T
- for some $\sigma \in Q^*$ we have

$$(q_{in}, 0) \Rightarrow^* (\sigma q', j)$$

$$(\sigma q', j) \Rightarrow^* (\sigma q' q, i)$$

and in the second computation $\sigma q'$ is never changed

Tabulation

1. $T = \{ (\perp, 0, q_{in}, 0) \}$
2. For $i = 1 \dots n$ do :
 - a. $N = \{ \}$
 - b. **shift** :
for each $(q', j, q_1, i-1)$ in T , $(q_1 \Rightarrow_{ai} q_1 q_2)$
in Δ , "add" $(q_1, i-1, q_2, i)$ to T and N
 - c. **switch** :
for each $(q_1, j, q_2, i-1)$ in T , $(q_1 q_2 \Rightarrow_{ai} q_1 q_3)$
in Δ , "add" (q_1, j, q_3, i) to T and N

ESSLLI 2004

Tabular Parsing

45

Tabulation

- d. while $N \neq \{ \}$ do :
 - i. remove some (q_1, j, q_2, i) from N
 - ii. **pop** :
for each (q', k, q_1, j) in T , $(q_1 q_2 \Rightarrow_{\varepsilon} q_3)$
in Δ , "add" (q', k, q_3, i) to T and N
3. Recognize iff $(\perp, 0, q_{fin}, n)$ in T

ESSLLI 2004

Tabular Parsing

46

Tabulation

Example (cont'd) :

- from $(\perp, 0, q_0, 0)$ in $\mathcal{T}$ and $q_0 \Rightarrow_a q_0 q_1$ in Δ with $a = a_1$, derive $(q_0, 0, q_1, 1)$ at step 2(b)
- from $(q_0, 0, q_1, 1)$ in $\mathcal{T}$ and $q_0 q_1 \Rightarrow_b q_0 q_2$ in Δ with $b = a_2$, derive $(q_0, 0, q_2, 2)$ at step 2(c)
- from $(q_0, 0, q_2, 2)$ and $(q_2, 2, q_6, 4)$ in $\mathcal{T}$, and $q_2 q_6 \Rightarrow_\varepsilon q_7$ in Δ , derive $(q_0, 0, q_7, 4)$ at step 2(d)

Tabulation

Implementation :

- use square **array** for $\mathcal{T}$
 - each element a **bit vector** of length $|Q|^2$
 - $T_{i,j}[q, q'] = \begin{cases} 1 & \text{if } (q', j, q, i) \in \mathcal{T} \\ 0 & \text{otherwise} \end{cases}$
- use **stack** or **heap** for $\mathcal{N}$

Tabulation

Time complexity dominated by step 2(d)ii :

- use square array for T
 - single execution time $O(1)$ (retrieve items, construct new item, store and check)
 - free variables are :
 - $(q_1 q_2 \Rightarrow_{\epsilon} q_3) \in \Delta$
 - $q' \in Q$
 - $i, j, k \in [0, \dots, n]$

Tabulation

Time complexity (cont'd) :

- total time is
$$O(|\Delta| |Q| n^3)$$
$$= O(|M| |Q| n^3)$$

Space complexity is $O(|Q|^2 n^2)$

Tabulation

1. $T = N = \{ (\perp, 0, q_{in}, 0) \}$
2. For $i = 0 \dots n$ do :
 - a. **shift** :
for each $(q', j, q_1, i-1)$ in T , $(q_1 \Rightarrow_{ai} q_1 q_2)$ in Δ , "add" $(q_1, i-1, q_2, i)$ to T and N
 - b. **α -switch** :
for each $(q_1, j, q_2, i-1)$ in T , $(q_1 q_2 \Rightarrow_{ai} q_1 q_3)$ in Δ , "add" (q_1, j, q_3, i) to T and N

ESSLLI 2004

Tabular Parsing

51

Tabulation

- c. while $N \neq \{ \}$ do :
 - i. remove some (q_1, j, q_2, i) from N
 - ii. **push** :
for each $(q_2 \Rightarrow_{\varepsilon} q_2 q_3)$ in Δ , "add" (q_2, i, q_3, i) to T and N
 - iii. **ε -switch** :
for each $(q_1 q_2 \Rightarrow_{\varepsilon} q_1 q_3)$ in Δ , "add" (q_1, j, q_3, i) to T and N

ESSLLI 2004

Tabular Parsing

52

Tabulation

iv. pop :

for each (q', k, q_1, j) in T , $(q_1 q_2 \Rightarrow_{\varepsilon} q_3)$
in Δ , "add" (q', k, q_3, i) to T and N ;
for each (q_2, i, q_3, i) in T , $(q_2 q_3 \Rightarrow_{\varepsilon} q')$
in Δ , "add" (q_1, j, q', i) to T and N

3. Recognize iff $(\perp, 0, q_{fin}, n)$ in T

Earley

Earley's algorithm :

- discovered in 1968
- strictly related to top-down parsing
- we specify the algorithm through tabularization and the definition of the [Earley automaton](#)
- we assume input CFG $G = (\Sigma, N, P, S)$ has a single production $S \rightarrow \sigma$ rewriting the start symbol

Earley

Earley automaton M_E :

- input alphabet is the same as G
- $Q = \{ [A \rightarrow \alpha a \bullet \alpha'] \mid A \rightarrow \alpha \alpha' \in P \}$;
elements of Q are called **dotted rules**
- $q_{in} = S \rightarrow \bullet \sigma$
- $q_{fin} = S \rightarrow \sigma \bullet$

Earley

M_E (cont'd) :

- Δ contains the following transitions :
 - $[A \rightarrow \alpha \bullet B \beta] \Rightarrow_\epsilon [A \rightarrow \alpha \bullet B \beta] [B \rightarrow \bullet \gamma]$,
for each $A \rightarrow \alpha B \beta$ and $B \rightarrow \gamma$, called **predict**
 - $[A \rightarrow \alpha \bullet b \beta] \Rightarrow_b [A \rightarrow \alpha b \bullet \beta]$,
for each $A \rightarrow \alpha b \beta$, called **shift**
 - $[A \rightarrow \alpha \bullet B \beta] [B \rightarrow \gamma \bullet] \Rightarrow_\epsilon [A \rightarrow \alpha B \bullet \beta]$,
for each $A \rightarrow \alpha B \beta$ and $B \rightarrow \gamma$, called **complete**

Earley

M_E (cont'd) :

- $|Q| = |G|$
- $|D| = O(|M|) = O(|G|^2)$
- M_E is nondeterministic
- predict, shift and complete correspond to transitions of the type push, $\bar{x}$ -switch and pop

Earley

M_E (cont'd) :

- Δ contains the following transitions :
 - $[A \rightarrow \alpha \bullet B \beta] \Rightarrow_\epsilon [A \rightarrow \alpha \bullet B \beta] [B \rightarrow \bullet \gamma]$,
for each $A \rightarrow \alpha B \beta$ and $B \rightarrow \gamma$, called **predict**
 - $[A \rightarrow \alpha \bullet b \beta] \Rightarrow_b [A \rightarrow \alpha b \bullet \beta]$,
for each $A \rightarrow \alpha b \beta$, called **shift**
 - $[A \rightarrow \alpha \bullet B \beta] [B \rightarrow \gamma \bullet] \Rightarrow_\epsilon [A \rightarrow \alpha B \bullet \beta]$,
for each $A \rightarrow \alpha B \beta$ and $B \rightarrow \gamma$, called **complete**

Earley

1. $T = N = \{ (0, [S \rightarrow \bullet \sigma], 0) \}$
2. For $i = 0 \dots n$ do :
 - a. **a-switch :**
for each $(j, [A \rightarrow \alpha \bullet a \alpha'], i-1)$ in T , $a = a_i$,
"add" $(j, [A \rightarrow \alpha a \bullet \alpha'], i)$ to T and N
 - b. while $N \neq \{ \}$ do :
 - i. remove some $(j, [A \rightarrow \alpha \bullet X \alpha'], i)$ from N

Earley

- ii. **push :**
if $X = B$, for each $B \rightarrow \beta$ in P , "add"
 $(i, [B \rightarrow \bullet \beta], i)$ to T and N
 - iii. **pop :**
if $X \alpha' = \varepsilon$, for each $(k, [B \rightarrow \beta \bullet A \beta'], j)$
in T , "add" $(k, [B \rightarrow \beta A \bullet \beta'], i)$ to T
and N ;
if $X = B$, for each $(i, [B \rightarrow \beta \bullet], i)$ in T ,
"add" $(j, [A \rightarrow \alpha B \bullet \alpha'], i)$ to T and N
3. Recognize iff $(0, q_{fin}, n)$ in T

Cocke-Kasami-Younger

1. $T = \{ \}$
2. For $i = 1 \dots n$ do :
 - a. **shift :**
for each $A \rightarrow a$ in P with $a = a_i$,
add $(i-1, A, i)$ to T
 - b. **pop :**
for each (k, B, j) and (j, C, i) in T ,
 $A \rightarrow BC$ in P , add (k, A, i) to T
3. Recognize iff $(0, S, n)$ in T

Generalized LR

1. $T = \{ (\perp, 0, q_{in}, 0) \}$
2. For $i = 1 \dots n$ do :
 - a. $N = \{ \}$
 - b. **shift :**
for each $(q', j, q_1, i-1)$ in T , $(q_1 \Rightarrow_a q_1 q_2)$
in Δ , $a = a_i$, "add" $(q_1, i-1, q_2, i)$ to T and N

Generalized LR

- c. while $N \neq \{ \}$ do :
 - i. remove some $(q_{m-1}, j_{m-1}, q_m, j_m)$ from N
 - ii. reduce :
 - for each $(q_0 q_1 \dots q_{m-1} q_m \Rightarrow_{\varepsilon} q_0 q')$ in Δ ,
 $(q_0, j_0, q_1, j_1), (q_1, j_1, q_2, j_2), \dots,$
 $(q_{m-2}, j_{m-2}, q_{m-1}, j_{m-1})$ in T ,
"add" (q_0, j_0, q', j_m) to T and N ;
 - for each $(q_{in} q_1 \dots q_{m-1} q_m \Rightarrow_{\varepsilon} q_{fin})$ in Δ ,
 $(\perp, 0, q_{in}, j_0), (q_{in}, j_0, q_1, j_1), \dots,$
 $(q_{m-2}, j_{m-2}, q_{m-1}, j_{m-1})$ in T ,
"add" $(\perp, 0, q_{fin}, j_m)$ to T and N

3. Recognize iff $(\perp, 0, q_{fin}, n)$ in T

ESSLLI 2004

Tabular Parsing

63