#### Lecture 10: Bottom-up parsing

David Hovemeyer

October 2, 2023

601.428/628 Compilers and Interpreters



► yacc and bison work by magic

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三三 - のへぐ

Any questions?

yacc and bison generate shift/reduce, bottom-up parsers

▲□▶ ▲□▶ ▲□▶ ▲□▶ ▲□ ● ● ●

- using the LALR(1) algorithm
- Today we'll investigate the basic principles
  - using a simpler algorithm called SLR

# Bottom-up parsing

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三三 - のへぐ

An approach to parsing that is more powerful (handles a larger set of possible grammars) than predictive (top-down) parsing.

The basic idea is to start with the string of terminal symbols, and use grammar productions to reduce the string to the grammar's start symbol. In other words, we're constructing a derivation in reverse.

Just + and \* operators, a and b as primary expressions

◆□▶ ◆□▶ ◆□▶ ◆□▶ □ のQ@

 $\begin{array}{l} \mathsf{E} \rightarrow \mathsf{E} + \mathsf{T} \\ \mathsf{E} \rightarrow \mathsf{T} \\ \mathsf{T} \rightarrow \mathsf{T} * \mathsf{F} \\ \mathsf{T} \rightarrow \mathsf{F} \\ \mathsf{F} \rightarrow \mathsf{a} \\ \mathsf{F} \rightarrow \mathsf{b} \end{array}$ 

In a leftmost derivation, each step replaces the leftmost nonterminal

In a rightmost derivation, each step replaces the rightmost nonterminal

A "backwards" derivation (each production "generates" the left-hand nonterminal by replacing the right-hand symbols)

Working string	Production
$\underline{a} + b * a$	$F \to a$
<u>F</u> + b * a	$T\toF$
$\underline{T} + b$ * a	$E\toT$
$E + \underline{b} * a$	$F \to b$
E + <u>F</u> * a	$T\toF$
E + T * <u>a</u>	$F \to a$
E + <u>T * F</u>	$T\toT*F$
E + T	$E\toE+T$
E	

Notice that when we read the series of steps in the parse from bottom to top, it forms a rightmost derivation of the input string.

A sentential form is a string of symbols (terminal and/or nonterminal) which can be derived from a grammar's start symbol.

a.k.a. "working string"

A right-sentential form is a sentential form that can be derived using a rightmost derivation from a grammar's start symbol.

For example, |E + T \* a| is a right-sentential form of the grammar above, because it can be derived from the start symbol E using a rightmost derivation.

As in top-down parsing, the main issue in bottom-up parsing is knowing when to apply a production. The notion of *handles* is helpful in making this determination.

Given a right sentential form  $\alpha\beta{\rm w},$  then  $\beta$  is a handle of the sentential form if the production

 $\mathsf{A}\to\beta$ 

is used in the rightmost derivation of the sentential form to expand

 $\alpha \mathsf{Aw}$ 

into

◆□▶ ◆□▶ ◆三▶ ◆三▶ ○ ◆○◇

If we can reliably find the handle  $\beta$  of any right-sentential form, then we know how to reduce  $\beta$  to make progress in the reverse construction of the rightmost derivation of the input string. This idea is called *handle pruning*.

Note that the string w (possibly empty) can only consist of terminal symbols. (If there were any nonterminal symbols in w, then reducing  $\beta$  would not be part of a reversed rightmost derivation.)

▲□▶ ▲□▶ ▲□▶ ▲□▶ ▲□ ● ● ●

Given the notion of handles and handle pruning, we can describe a simple approach, *shift/reduce parsing*, for bottom-up parsing.

Shift-reduce parsing uses two data structures:

- the input string, followed by the special \$ terminal symbol to mark the end of the string
- ▶ a stack of symbols, initially containing the special \$ symbol

At each step, a shift/reduce parser has two choices for what to do next:

- 1. Shift the next terminal symbol from the input string onto the stack
- 2. Reduce the handle which appears on the top of the stack

#### Example shift/reduce parse



#### Example shift/reduce parse

Using input string |a + b \* a|Stack Iuput string Achon \$ a + 6 \* a \$ shift + 6 \* a\$ reduce F > a \$ E T reduce T > F ł 6 <del>\*</del> a \$ b \* a \$ reduce E > T \$ E +b \* a \$ shift \$E+ \* a \$ shift \* reduce F > 6 \$ E + ba \$ reduce T > F E+ P a \$ \* \* E+T 9 \$ shift \$ E + T\* 2 shift E+T\*a reduce F > a \$ E + T \* F reduce TAT\*F reduce E>E+T \$ E + T Ż \$E accept

## Parser construction

So, how can we automatically construct a bottom-up parser?

From the grammar, derive a set of *items*. An item represents, for a particular production, where the parser might be with respect to (eventually) reducing the production.

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

For example, in the production  $E \rightarrow E + T$ 

there are 4 items:

$$E \rightarrow \bullet E + T$$
$$E \rightarrow E \bullet + T$$
$$E \rightarrow E + \bullet T$$
$$E \rightarrow E + \bullet T$$
$$E \rightarrow E + T \bullet$$

The dot  $(\bullet)$  indicates how close the parser is to reducing the production. In the case of the first item,

 $\mathsf{E} \to \bullet \; \mathsf{E} + \mathsf{T}$ 

the parser "wants" to apply the production, and is about to start working parsing the initial E nonterminal on the right-hand side of the production.

In the last item

 $\mathsf{E}\to\mathsf{E}+\mathsf{T}\bullet$ 

the parser has successfully parsed the entire right hand side of the production, and can consider reducing the production. (In other words, it is possible that the handle E + T is now on the top of the stack.)

A challenge of bottom-up parsing is that at any given point in parsing, many items may be "active" simultaneously. So, the parser uses the notion of *item sets* to represent this possibility.

The LR(0) automaton is an automaton we can construct to guide a bottom-up parser in making parsing decisions.

We start by constructing an augmented grammar. The augmented grammar adds a single production,

 $S'\!\!\to S$ 

where S is the original grammar's start symbol. The parser will know that it has completed parsing when the original start symbol S is reduced to the augmented grammar's start symbol S'.

Next, we construct the canonical LR(0) item sets. To do so, we need to define two functions, CLOSURE and GOTO.

The CLOSURE function builds a "complete" item set starting from an initial item set. CLOSURE(I), where I is an item set, is defined as follows:

Repeatedly: If  $A \to \alpha \bullet B\beta$  is an item in I, then add all items  $B \to \bullet \gamma$  to CLOSURE(I).

The item sets constructed by the CLOSURE function form the states of the LR(0) automaton.

The GOTO function is used the construct the transitions of the LR(0) automaton. GOTO(I, X), where I is an item set and X is a grammar symbol (terminal or nonterminal), is defined as follows

If  $[A \rightarrow \alpha \bullet B\beta]$  is an item in I, then the closure of all items  $[A \rightarrow \alpha B \bullet \beta]$  is in GOTO(I, B)

Basically, the GOTO function describes how the dot  $(\bullet)$  is moved when a nonterminal is reduced.

◆□▶ ◆□▶ ◆三▶ ◆三▶ 三三 のへぐ

C is the set of all LR(0) item sets. To start, put CLOSURE( $\{S' \rightarrow \bullet S\}$ ) in C.

▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

Repeat:

For each set I in C For each grammar symbol X if GOTO(I,X) is nonempty and is not already in C add GOTO(I, X) to C add a transition from I to GOTO(I, X) on symbol X

until no more item sets are added to C

## Starting the LR(0) automaton

 $\begin{array}{l} \mathsf{E}' \rightarrow \mathsf{E} \\ \mathsf{E} \rightarrow \mathsf{E} + \mathsf{T} \\ \mathsf{E} \rightarrow \mathsf{T} \\ \mathsf{T} \rightarrow \mathsf{T} * \mathsf{F} \\ \mathsf{T} \rightarrow \mathsf{F} \\ \mathsf{F} \rightarrow \mathsf{a} \\ \mathsf{F} \rightarrow \mathsf{b} \end{array}$ 

・ロト・日本・日本・日本・日本・日本

#### LR(0) construction example

Part of the LR(0) automaton for our expression grammar, showing the initial item set  $I_0$  and the item sets added by computing the GOTO sets from the initial item set



#### Complete LR(0) automaton

To complete the automaton, we continue the algorithm, adding additional states and transitions as needed.



# SLR parsing

◆□ > ◆□ > ◆ 三 > ◆ 三 > ● ○ ○ ○ ○

SLR ("Simple LR") is a simple bottom-up parsing algorithm that makes use of the LR(0) automaton.

Basic idea: keep a stack of states in addition to the input string and symbol stack. Each symbol on the symbol stack has a matching state on the state stack.

The top state tells us where we (currently) are in the LR(0) automaton.

The parser must decide, at each step, whether to shift a terminal symbol onto the stack or reduce a production. It uses the LR(0) automaton to guide this decision (explained on next slide.)

If the current (top) state has a transition on the next terminal symbol in the input string,

- shift it onto the symbol stack
- whatever state is reached by following the transition, push the state on the state stack

Otherwise, reduce. One of the items in the current state will be of the form

 $\mathsf{A} \to \alpha \bullet$ 

indicating that the symbols matching  $\alpha$  are cleared from the symbol stack and A pushed on. The states corresponding to  $\alpha$  are cleared from the state stack. After clearing the states, we return to whatever state is currently on the top of the state stack and look for a transition on A. Whatever state the transition leads to is pushed on the state stack.

## Example SLR parse

State Stack	Symbol Stack	Input	Action	
\$ 0	\$	a + b * a \$		

## Example SLR parse

State Stack	Symbol Stack	Input	Action
\$ 0	\$	a + b * a \$	shift
\$04	\$ a	+ b * a \$	$reduce\;F\toa$
\$03	\$ F	+ b * a \$	$reduce\;T\toF$
\$02	\$ T	+ b * a \$	$reduce\;E\toT$
\$01	\$ E	+ b * a \$	shift
\$016	\$E+	b*a\$	shift
\$0165	\$ E + b	* a \$	$reduce\;F\tob$
\$0163	\$ E + F	* a \$	$reduce\;T\toF$
\$0168	\$ E + T	* a \$	shift
\$01687	\$ E + T *	а\$	shift
\$016874	\$ E + T * a	\$	$reduce\;F\toa$
\$016879	\$ E + T * F	\$	reduce T $\rightarrow$ T * F
\$0168	\$ E + T	\$	$reduce\;E\toE+T$
\$01	\$ E	\$	accept

## Table-driven bottom-up parsing

◆□▶ ◆□▶ ◆ □▶ ◆ □▶ ○ □ ○ ○ ○ ○

There is a general way to construct a shift-reduce parser. The information about how to decide which action (shift/reduce/accept/error) the parser should take at any point is encoded in an *LR parse table*.

An LR parse table consists of columns for each terminal and nonterminal symbol. The terminal columns are the ACTION columns, while the nonterminal columns are the GOTO columns.

ACTION columns allow the parser to make a decision, based on the next terminal symbol in the input string, what action to take. GOTO columns determine how to determine which state to go to next when a production is reduced.

ACTION[i, a] is the ACTION entry for state i when terminal a is seen. GOTO[i, A] is the state to go to when nonterminal A is reduced in state i. Construct the LR(0) automaton for the grammar. Number the productions of the grammar.

For each transition from state i to k in the LR(0) automaton labeled with a terminal symbol a, set ACTION[i, a] to **s**k, meaning "shift state k onto the stack".

For each transition from state i to k in the LR(0) automaton labeled with a nonterminal symbol A, set GOTO[i, A] to k. This means that if we reduce a handle on the stack to an A nonterminal, clearing the handle from the stack and leaving state i on top of the stack, we should push state k.

For each state i with an item of the form  $A \rightarrow \alpha \bullet$ , where A is any nonterminal except S' (the start symbol of the augmented grammar), then for each terminal symbol a in FOLLOW(A), set ACTION[i, a] to **r**k, meaning "reduce production k", where k is the number of the production  $A \rightarrow \alpha$ .

If state i contains the item S'  $\to$  S  $\bullet,$  then set ACTION[i, \$] to acc, meaning "accept".

▲□▶ ▲□▶ ▲□▶ ▲□▶ ■ ●の00

The resulting parse table is called the SLR(1) parse table.

◆□ > ◆□ > ◆ 三 > ◆ 三 > ● ○ ○ ○ ○

Grammar productions (numbered):

$$\begin{array}{lll} (1) & \mathsf{E}' \to \mathsf{E} \\ (2) & \mathsf{E} \to \mathsf{E} + \mathsf{T} \\ (3) & \mathsf{E} \to \mathsf{T} \\ (4) & \mathsf{T} \to \mathsf{T} & \mathsf{F} \\ (5) & \mathsf{T} \to \mathsf{F} \\ (5) & \mathsf{F} \to \mathsf{a} \\ (6) & \mathsf{F} \to \mathsf{a} \\ (7) & \mathsf{F} \to \mathsf{b} \end{array}$$

Nonterminal	FIRST set	FOLLOW set
E'	{a, b}	{ \$ }
Е	{a, b}	{ +, \$ }
Т	{a, b}	{ *, +, \$ }
F	{a, b}	{ *, +, \$ }

#### Constructed SLR(1) parse table

Based on LR(0) automaton shown previously; blank entries are *error* actions

	ACTION					G	ГО	0
State	а	b	+	*	\$	Ε	Т	F
0	s4	s5				1	2	3
1			s6		acc			
2			r3	s7	r3			
3			r5	r5	r5			
4			r6	r6	r6			
5			r7	r7	r7			
6	s4	s5					8	3
7	s4	s5						9
8			r2	s7	r2			
9			r4	r4	r4			

Given an LR parse table, we can automate parsing of an input string.

Note that the general LR parser uses a stack of states rather than a stack of symbols; however, each state (except for state 0) is uniquely associated with a particular symbol, because all incoming edges to a particular state must be on the same symbol. (This is a trivial consequence of the way in which the LR(0) automaton is constructed; the transitions are generated by the computation of GOTO sets on specific symbols.)

#### Example table-driven parse

Step	Stack	Input	Action
1	\$ 0	a + b * a \$	

## Example SLR(1) table-driven parse

Step	Stack	Input	Action
1	\$ 0	a + b * a \$	shift 4
2	\$04	+ b * a \$	reduce 6 (F $ ightarrow$ a), goto 3
3	\$03	+ b * a \$	reduce 5 (T $ ightarrow$ F), goto 2
4	\$02	+ b * a \$	reduce 3 (E $ ightarrow$ T), goto 1
5	\$01	+ b * a \$	shift 6
6	\$016	b*a\$	shift 5
7	\$0165	* a \$	reduce 7 (F $ ightarrow$ b), goto 3
8	\$0163	* a \$	reduce 5 (T $\rightarrow$ F), goto 8
9	\$0168	* a \$	shift 7
10	\$01687	a \$	shift 4
11	\$016874	\$	reduce 6 (F $ ightarrow$ a), goto 9
12	\$016879	\$	reduce 4 (T $\rightarrow$ T * F), goto 8
13	\$0168	\$	reduce 2 (E $ ightarrow$ E $+$ T), goto 1
14	\$01	\$	accept

# Conflicts

▲□▶ ▲□▶ ▲ □▶ ▲ □▶ ▲ □ ● ● ● ●

Sometimes when we construct the parse tables we'll encounter a situation where we have conflicting actions in an entry. This may indicate that the grammar is ambiguous, or it could indicate that the parsing algorithm simply isn't powerful enough to handle the grammar.

Shift/reduce conflict: we don't know whether to shift a token or reduce a production.

Reduce/reduce conflict: there are two possible productions which we could reduce.

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

#### Example shift/reduce conflict



The problem is that the terminal symbol c is in the FOLLOW set of A (because of the S  $\rightarrow$  b A c production), so when the parser sees a "c" in state 1, it doesn't know whether to shift or reduce.



▲ロ ▶ ▲周 ▶ ▲ 国 ▶ ▲ 国 ▶ ● の Q @

In state 2, we won't know which production to reduce.

Idea: with each item, keep track of which terminal symbols could follow immediately after a reduction of the production. This is called an LR(1) item.

Example LR(1) item:

 $\mathsf{A} \to \alpha \bullet$ , a

means that a state containing the item, the production  $A \rightarrow \alpha$  is reduced if we see an a as the next input symbol.

Keeping track of this additional information gives the parser more context to decide when a production should be reduced.

There are more powerful ways to construct LR parse tables, based on LR(1) items:

- LR: powerful, but generates very large parse tables
- LALR: nearly as powerful as LR (at least for grammars that typically arise in the implementation of a programming language), but the parse tables are generally as small as SLR(1) tables

Most bottom-up parser generators (e.g., yacc, bison) use the LALR algorithm to construct the parse table.

▲□▶ ▲□▶ ▲□▶ ▲□▶ ■ ●の00