

CS4613 Lecture 3: Evaluation and AE

David Bremner

March 21, 2023

Semantics (= Evaluation)

We can use BNF specifications to help define **meaning**. Given a grammar

```
ae := NUMBER
    | ae1 + ae2
    | ae1 - ae2
```

```
ae1: ae
```

```
ae2: ae
```

We can define an **evaluator** function (pseudocode)

- `eval(NUMBER) = NUMBER(?)`
- `eval(ae1 + ae2) =
eval(ae1) + eval(ae2)`
- `eval(ae1 - ae2) =
eval(ae1) - eval(ae2)`

└ Evaluation as Semantics

└ Semantics (= Evaluation)

Semantics (= Evaluation)

We can use BNF specifications to help define **meaning**. Given a grammar

```
ae := NUMBER
    | ae1 + ae2
    | ae1 - ae2
ae1: ae
ae2: ae
```

We can define an **evaluator** function (pseudocode)

```
a. eval(NUMBER) = NUMBER(?)
b. eval(ae1 + ae2) =
  eval(ae1) + eval(ae2)
c. eval(ae1 - ae2) =
  eval(ae1) - eval(ae2)
```

1. meaning in our context is "running" a program (or "interpreting", "compiling", but we will use "evaluating").
2. The text (and many others) use the convention of $\langle A \rangle$ to denote meta-variables, i.e. variables whose values are pieces of program.

Note the completely different roles of the two “+”s and “-”s. In fact, it might have been more correct to write:

a. `eval(NUMBER) = NUMBER`

b. `eval(ae1 "+" ae2) = eval(ae1) + eval(ae2)`

c. `eval(ae1 "-" ae2) = eval(ae1) - eval(ae2)`

When we write:

`eval(1 - 2 + 3) = ... = 1 - 2 + 3`

we have two expressions, but one stands for an input syntax, and one stands for a ‘real’ mathematical expression.

└ Evaluation as Semantics

Note the completely different roles of the two "+"s and "-".s. In fact, it might have been more correct to write:

a. `eval(NUMBER) = NUMBER`
 b. `eval(ae1 "+" ae2) = eval(ae1) + eval(ae2)`
 c. `eval(ae1 "-" ae2) = eval(ae1) - eval(ae2)`

When we write:

`eval(1 - 2 + 3) = ... = 1 - 2 + 3`

we have two expressions, but one stands for an input syntax, and one stands for a 'real' mathematical expression.

- Note that there's a similar kind of informality in our **syntax** BNF specifications, where we assume that FOO (foo) refers to some terminal (non-terminal).
- An alternative popular notation for `eval(X)` is `[[X]]`:
 - `[[NUMBER]] = NUMBER`
 - `[[ae1 + ae2]] = [[ae1]] + [[ae2]]`
 - `[[ae1 - ae2]] = [[ae1]] - [[ae2]]`

Ambiguity

`eval(1 - 2 + 3) = ?`

Depending on the way the expression is parsed, 2 or -4:

```
eval(1 - 2 + 3) = eval(1 - 2) + eval(3)
                  ; [b]
                  = eval(1) - eval(2) + eval(3)
                  ; [c]
                  = 1 - 2 + 3
                  ; [a,a,a]
                  = 2
```

```
eval(1 - 2 + 3) = eval(1) - eval(2 + 3)
                  ; [c]
                  = eval(1) - (eval(2) + eval(3))
                  ; [a]
                  = 1 - (2 + 3)           ; [a,a,a]
                  = -4
```

- ▶ We need parens around a sub-expression only in one case, why?
- ▶ With 'eval' the parse tree matters, so ambiguity must be eliminated.

└ Evaluation as Semantics

```
eval(1 - 2 + 3) = eval(1) - eval(2 + 3) ; [c]
                = eval(1) - (eval(2) + eval(3)) ; [c]
                = 1 - (2 + 3) ; [a]
                = -4 ; [a,a,a]
```

- ▶ We need parens around a sub-expression only in one case, why?
- ▶ With 'eval' the parse tree matters, so ambiguity must be eliminated.

1. In a case of a computer implementation, the syntax on the left is (as always) an AE syntax, and the 'real' expression on the right is an expression in whatever language we use to implement our AE language.

Suppose we want to parse numbers ourselves

number: DIGIT number

`eval(0) = 0`

`eval(1) = 1`

...

`eval(9) = 9`

`eval(DIGIT) = DIGIT`

`eval(DIGIT number) = 10*eval(DIGIT) + eval(number)`

This grammar is unambiguous.

number: DIGIT | DIGIT number

But the corresponding eval rule is wrong:

$$\begin{aligned} \text{eval}(123) &= 10 * \text{eval}(1) + \text{eval}(23) \\ &= 10 * 1 + 10 * \text{eval}(2) + \text{eval}(3) \\ &= 10 * 1 + 10 * 2 + 3 \\ &= 33 \end{aligned}$$

Changing the order of the last rule works much better:

```
number: DIGIT | number DIGIT
```

And then:

```
eval(number DIGIT) = 10*eval(number) + eval(DIGIT)
```

- ▶ What are potential problems (think implementation language versus implemented language).

└ Evaluation as Semantics

Changing the order of the last rule works much better:

```
number: DIGIT | number DIGIT
```

And then:

```
eval(number DIGIT) = 10*eval(number) + eval(DIGIT)
```

- ▶ What are potential problems (think implementation language versus implemented language).

1. Think of an average language that does not give you bignums, making the above rules fail when the numbers are too big. In Racket, we happen to be using an integer representation for the syntax of integers, and both are unlimited. But what if we wanted to write a Racket compiler in C or a C compiler in Racket? What about a C compiler in C, where the compiler runs on a 64 bit machine, and the result needs to run on a 32 bit machine?

Compositionality

We saw this grammar is easier to evaluate:

```
number: DIGIT | number DIGIT
```

- ▶ The corresponding language is **compositional**; the meaning of an expression is composed of the meaning of its parts (in a simple way).

└ Evaluation as Semantics

└ Compositionality

We saw this grammar is easier to evaluate:

number: DIGIT | number DIGIT

- ▶ The corresponding language is **compositional**; the meaning of an expression is composed of the meaning of its parts (in a simple way).

1. In this case this can be tolerable, since the meaning of the expression is still made out of its parts – but imperative programming (when you use side effects) is much more problematic since it is not compositional (at least not in the obvious sense).
2. This is compared to functional programming, where the meaning of an expression is a combination of the meanings of its subexpressions. For example, every sub-expression in a functional program has some known meaning, and these all make up the meaning of the expression that contains them – but in an imperative program we can have a part of the code be 'x++' – and that doesn't have a meaning by itself, at least not one that contributes to the meaning of the whole program in a direct way.

The 'DIGIT number' case is more complicated. The meaning depends also on the input **syntax**

$$\text{eval}(\text{DIGIT number}) = \text{eval}(\text{DIGIT}) * 10^{\text{length}(\text{number})} + \text{eval}(\text{number})$$

Implementing an Evaluator

For simple languages, a merged parser and evaluator is possible

```
14 (define (eval sx)
    (let ([rec (lambda (fn)
                 (eval (fn (s-exp->list sx)))))]
      (cond
        [(s-exp-match? `NUMBER sx)
         (s-exp->number sx)]
        [(s-exp-match? `(+ ANY ANY) sx)
         (+ (rec second) (rec third))]
        [(s-exp-match? `(- ANY ANY) sx)
         (- (rec second) (rec third))]
        [else (error 'eval (to-string sx))]))))
```


Abstract and Concrete Syntax

- ▶ For more complex languages, evaluating in the parser is not usually practical.
- ▶ So we split into (at least) two layers
- ▶ By using the abstract syntax tree (AST) as an API, we can independently change the (concrete) syntax and the semantics.
- ▶ We can also e.g. provide better error checking/messages.

15 (eval `{+ 1 {- 3 "a"}})

└ Implementing an Evaluator

└ Abstract and Concrete Syntax

- ▶ For more complex languages, evaluating in the parser is not usually practical.
- ▶ So we split into (at least) two layers
- ▶ By using the abstract syntax tree (AST) as an API, we can independently change the (concrete) syntax and the semantics.
- ▶ We can also e.g. provide better error checking/messages.

```
▣ (eval '(+ 1 {- 3 "a"}))
```

1. We'll see that typechecking actually also works on the abstract syntax.
2. this is like the distinction between XML syntax and well-formed XML syntax.

Implementing The AE Language

Back to our 'eval' – this will be its (obvious) type:

```
(AE -> Number)
;; consumes an AE and computes the corresponding number
```

Which leads to some obvious test cases:

```
(test (eval (parse-sx `3))           3)
(test (eval (parse-sx `{+ 3 4}))     7)
(test (eval (parse-sx `{+ {- 3 4} 7})) 6)
```

We're usually less interested in testing the parser (except perhaps for debugging).

18

```
(test (parse-sx `{+ {- 3 4} 7})
      (Add (Sub (Num 3) (Num 4)) (Num 7)))
```

The structure of the recursive 'eval' code follows the recursive structure of its input.

```
19 (define (eval expr)
    (type-case AE expr
      [(Num n)      ....]
      [(Add l r)   (... (eval l) ... (eval r))]
      [(Sub l r)   (... (eval l) ... (eval r))]))
```

In this case, filling in the gaps is very simple

```
20 (define (eval expr)
    (type-case AE expr
      [(Num n)      n]
      [(Add l r)   (+ (eval l) (eval r))]
      [(Sub l r)   (- (eval l) (eval r))]))
```

We can further combine 'eval' and 'parse' into a single 'run' function that evaluates an AE s-expr.

```
(define (run sx)
  (eval (parse-sx sx)))
```

Four functions, just like grandpas calculator

```
#| BNF for the AE language:
  <AE> ::= <num>
        | { + <AE> <AE> }
        | { - <AE> <AE> }
        | { * <AE> <AE> }
        | { / <AE> <AE> }

|#
;; AE abstract syntax trees
(define-type AE
  [Num (val : Number)]
  [Add (l : AE) (r : AE)]
  [Sub (l : AE) (r : AE)]
  [Mul (l : AE) (r : AE)]
  [Div (l : AE) (r : AE)])
```

Putting the pieces together

```
25 (define (eval expr)
    (type-case AE expr
      [(Num n)      n]
      [(Add l r)    (+ (eval l) (eval r))]
      [(Sub l r)    (- (eval l) (eval r))]
      [(Mul l r)    (* (eval l) (eval r))]
      [(Div l r)    (/ (eval l) (eval r))]))

(define (run sx)
  (eval (parse-sx sx)))

(test (run `3)      3)
(test (run `{+ 3 4}) 7)
(test (run `{+ {- 3 4} 7}) 6)
```