

Lecture 4: Inference Rules

Steven Holtzen

s.holtzen@northeastern.edu

CS4400/5400 Fall 2024

1 Inference Rules

- Inference rules are powerful notation for describing recursive algorithms and inductively-defined data-structures.
- We will use them to describe interpreters, type systems, and prove things about programs.
- We'll start by revisiting a recursive function we saw earlier, the factorial function:

$$\text{fact}(n) = \begin{cases} 1 & \text{if } x = 0 \\ n \times \text{fact}(n-1) & \text{otherwise.} \end{cases} \quad (1)$$

- The above uses standard “math notation” you may have seen before in another class. Here is another way of writing the same program using Racket:

```
> (define (fact n)
  (if (equal? n 0) 1 (× n (fact (- n 1)))))
> (fact 4)
24
```

- There is a *third* style of notation for describing the `fact` function called **inference rules**, which look like this:

$$\frac{}{\text{fact}(0) = 1} \text{ (FACTBASE)} \qquad \frac{\text{fact}(n-1) = v}{\text{fact}(n) = n \times v} \text{ (FACTIND)}$$

- Inference rules consist of two pieces: **premises** and **conclusions**. The premises go above the horizontal line, and the conclusions go below the horizontal line.
- An inference rule with no premises is called an **axiom**. The FACTBASE rule above is an axiom. Axioms are always true; these are the base cases of your induction.
- The FACTIND rule has both premises and conclusions. We read it as: “if $\text{fact}(n-1) = v$ for some value v , then $\text{fact}(n) = n \times v$.”

2 Proving things using inference rules

- We can use these inference rules prove that $\text{fact}(3) = 6$ by drawing a **derivation tree**.
- The process works like this: first, we copy and paste what we want to prove as a conclusion to some to-be-determined inference rule:

$$\frac{???}{\text{fact}(3) = 6} \quad (2)$$

- Now we ask: *which inference rule can we apply to show this desired conclusion?* There is only one possible rule, `FactInd`:

$$\frac{\text{fact}(2) = v}{\text{fact}(3) = 3 \times v} \text{ (FACTIND)} \quad (3)$$

- However, we are not done: we haven't yet shown what $\text{fact}(2)$ is equal to. We can only terminate our proof when we hit a base case (i.e., an axiom). So, we keep going and apply the `FactInd` rule again:

$$\frac{\frac{\text{fact}(1) = v_1}{\text{fact}(2) = 2 \times v_2} \text{ (FACTIND)}}{\text{fact}(3) = 3 \times v} \text{ (FACTIND)} \quad (4)$$

- We are still not done, there is one more step of derivation:

$$\frac{\frac{\frac{\text{fact}(0) = 1 \text{ (FACTBASE)}}{\text{fact}(1) = 1 \times 1} \text{ (FACTIND)}}{\text{fact}(2) = 2 \times 1} \text{ (FACTIND)}}{\text{fact}(3) = 3 \times 2 \times 1} \text{ (FACTIND)} \quad (5)$$

- Now we are done. The above tree is called a derivation tree, and it is typically drawn bottom-up. Each leaf (top of the tree) must be an axiom.
- Drawing this derivation tree is essentially *running a factorial program by hand*. Each recursive case (an instance of `FACTIND`) requires performing a recursive call, which is either another inductive case or a base case. We label the result of the recursive call with a variable (like v). Once we hit a base case, we can replace the variables with concrete computed values. This is why in Equation (5) we substituted computed quantities (numbers) for variables (v, v_1, \dots).
- Once again we are in a host semantics situation: how do we interpret the meaning of numbers and operations like \times in derivation trees like those in Equation (5)? Once again we will use Racket as our host semantics: we will assume that these numbers and operations represent Racket numbers and Racket operations.

3 Inference Rules for a Calculator Interpreter

- Now let's use our new inference rule notation to describe a calculator interpreter
- We will define a function `eval` that evaluates each syntactic calculator term to a number
- Recall the syntax for a calculator language:

```
(struct (enum n))
(struct (eadd e1 e2))
(struct (emul e1 e2))
```

- We can define inference rules that describe the semantics of our calculator language:

$$\frac{}{(\text{eval } (\text{enum } n)) = n} \text{ (E-NUM)} \quad \frac{(\text{eval } e1) = v_1 \quad (\text{eval } e2) = v_2}{(\text{eval } (\text{eadd } e1 \ e2)) = (+ v_1 \ v_2)} \text{ (E-ADD)}$$

$$\frac{(\text{eval } e1) = v_1 \quad (\text{eval } e2) = v_2}{(\text{eval } (\text{emul } e1 \ e2)) = (* v_1 \ v_2)} \text{ (E-MUL)}$$

- Now we can draw derivation trees that describe evaluating our interpreter. For example, we can run the program `(+ (+ 1 2) 3)`:

$$\frac{\frac{(\text{eval } 1) = 1 \quad (\text{eval } 2) = 1}{(\text{eval } (+ \ 1 \ 2)) = 3} \text{ (E-ADD)} \quad \frac{}{(\text{eval } (3)) = 3} \text{ (E-NUM)}}{(\text{eval } (+ \ (+ \ 1 \ 2) \ 3)) = 6} \text{ (E-ADD)}$$

Note: we will sometimes use surface syntax instead of abstract syntax, and elide the name of rules, to make drawing derivation trees more concise.

- Notice how these rules visualize running the interpreter. For this reason, we will sometimes refer to drawing a derivation tree for a program as **running an interpreter by hand**.

4 Substitution Semantics for Let Language

- The “evaluates to” inductive definition is so common that we will introduce special notation to describe it: the relation $e \Downarrow v$ says e evaluates to v . This means exactly the same thing as $\text{eval}(e) = v$; it’s just easier (and common convention) to write it this way. This is sometimes called the *natural semantics* of a programming language.
- Let’s use this notation to describe the semantics for the `let` language from last lecture:

$$\frac{}{(enum\ n) \Downarrow n} \text{ (E-NUM)} \quad \frac{e_1 \Downarrow v_1 \quad e_2 \Downarrow v_2}{(eadd\ e_1\ e_2) \Downarrow (+\ v_1\ v_2)} \text{ (E-ADD)}$$

$$\frac{e_1 \Downarrow v_1 \quad e_2[x \mapsto v_1] \Downarrow v_2}{(\text{let}\ x\ e_1\ e_2) \Downarrow v_2} \text{ (E-LET)}$$

- Recall that the syntax $e_2[x \mapsto v_1]$ substitutes the value v_1 for x in e_2 .
- Let’s look at the rule for E-LET closely. What it says is: if e_1 evaluates to a value v_1 and $e_2[x \mapsto v_1]$ evaluates to v_2 , then the expression $(\text{let}\ x\ e_1\ e_2)$ evaluates to v_2 . Notice how there is a data dependence between the two premises: the value v_1 is referred to by both the rule for evaluating e_1 and e_2 . This is permitted: we just need to ensure that this value v_1 is *the same* in both locations.
- Now we can run a `let` program by hand. Let’s run the program $(\text{let}\ x\ 10\ (+\ x\ 5))$:

$$\frac{\frac{}{10 \Downarrow 10} \text{ (E-NUM)} \quad \frac{10 \Downarrow 10 \quad 5 \Downarrow 5}{(+\ x\ 5)[x \mapsto 10] \Downarrow 15} \text{ (E-ADD)}}{(\text{let}\ x\ 10\ (+\ x\ 5)) \Downarrow 15} \text{ (E-LET)}$$

- We will handle errors differently in inference rule notation than in Racket. In our inference rule notation, it is considered an error if we get “stuck” (meaning, there is no rule to apply to make progress).
- For example, if we try to run the program that consists solely of an identifier x , then we get stuck immediately because there is no rule to evaluate an identifier.

5 Environment-passing Semantics for Let Language

- The problem with substitution is that it is inefficient: we wouldn't want to actually use substitution to implement a programming language.
- A more practical solution to implementing the `let` language is to use an **environment** that maps identifiers to values. The environment will store the current value for a particular variable, and when an identifier is encountered during execution its value can be looked up in the environment.
- We will use an immutable hash table that maps strings (identifiers) to numbers as our environment. Immutable hash tables have the follow usage in Racket:

```
> (define my-tbl (hash)) ; declare an immutable empty hash-table
> (define my-tbl-with-x (hash-set my-tbl "x" 10)) ; add "x" with value 10
> (hash-has-key? my-tbl-with-x "x")
#t
> (hash-has-key? my-tbl "x") ; check if my-tbl contains "x"
#f
> (hash-ref my-tbl-with-x "x") ; look up "x" in my-tbl-with-x
10
```

- Here is an example implementation in Racket of an environment-passing interpreter:

```
;;; type expr =
;;; | add of expr × expr
;;; | mul of expr × expr
;;; | num of number
;;; | elet of string × expr × expr
(struct eadd (e1 e2) #:transparent)
(struct enum (n) #:transparent)
(struct elet (id assignment body) #:transparent)
(struct eident (id) #:transparent)

;;; eval : expr → (string, number) hashtable → number
;;; evaluates an expression e in environment env to a number
(define (eval e env)
  (match e
    [(enum n) n]
    [(eadd e1 e2) (+ (eval e1 env) (eval e2 env))]
    [(elet id assignment body)
     (let× [(v-assgn (eval assignment env))
            (new-env (hash-set env id v-assgn))]
       (eval body new-env))]
    [(eident id)
     (if (hash-has-key? env id)
         (hash-ref env id)
         (error "unbound identifier"))]))

(check-equal? (eval (elet "x" (enum 10) (eident "x"))) (hash)) 10)
(check-equal? (eval (elet "x" (enum 10)
                      (elet "x" (enum 20) (eident "x")))) (hash)) 20)
```

- **Check:** Why does the above implementation implement lexical scope? What property of the hash tables ensures this?

6 Inference Rules for Environment-passing Semantics of Let Language

- Let's get more practice with inference rules by using them to describe the environment-passing interpreter for the `let` language
- We will use the symbol ρ (pronounced "rho") for our environment. Looking up an identifier x in ρ is written $\rho[x]$ (this corresponds to `hash-ref`). Immutably updating the table ρ to set x equal to v is denoted $\rho[x \mapsto v]$ (this corresponds to `hash-set`). Checking if ρ contains x is denoted $x \in \rho$ (this corresponds with `hash-has-key?`).
- Now we can write down the environment-passing semantics of the `let` language in terms of inference rules. These rules will define a function $\langle e, \rho \rangle \Downarrow v$, which is read "expression e with environment ρ evaluates to value v ":

$$\frac{}{\langle (\text{enum } n), \rho \rangle \Downarrow n} \text{ (E-NUM)} \quad \frac{e_1 \Downarrow v_1 \quad e_2 \Downarrow v_2}{\langle (\text{eadd } e_1 \ e_2), \rho \rangle \Downarrow (+ \ v_1 \ v_2)} \text{ (E-ADD)}$$

$$\frac{x \in \rho \quad \rho[x] = v}{\langle x, \rho \rangle \Downarrow v} \text{ (E-IDENT)}$$

$$\frac{\langle e_1, \rho \rangle \Downarrow v_1 \quad \langle e_2, \rho[x \mapsto v_1] \rangle \Downarrow v_2}{\langle (\text{let } x \ e_1 \ e_2), \rho \rangle \Downarrow v_2} \text{ (E-LET)}$$

- **Notice:** the above rules differ quite substantially from the substitution-based semantics for `let`. One big difference is that we now have a rule for evaluating identifiers.
- Now we can draw derivation trees that visualize running our environment-passing interpreter.
- **Example:** Draw the derivation tree that corresponds to evaluating `(let "x" 10 x)` with the environment initially empty ($\rho = \{\}$):

$$\frac{\frac{10 \Downarrow 10 \text{ (E-NUM)} \quad \frac{x \in \{x \mapsto 10\} \quad \{x \mapsto 10\}[x] = 10}{\langle x, \{x \mapsto 10\} \rangle \Downarrow 10} \text{ (E-IDENT)}}{\langle (\text{let } "x" \ 10 \ x), \{\} \rangle \Downarrow 10} \text{ (E-LET)}}$$