

CPSC 311: Definition of Programming Languages: Environments

18-env

DRAFT

Joshua Dunfield
University of British Columbia

November 2, 2015

1 Midterm question titles

- “A substitutable cat”: I haven’t seen a cat near SWNG, but I have seen a fox on the West Mall sidewalk.
- “The Future”: Various songs have this title; the one I was thinking of is by the Toronto-based band Austra.
- “Add All Things”: A room in one of the McGill libraries has “PROVE ALL THINGS” carved into the wall.
- “Little Perennials”: a song by the Indigo Girls.
- “And so it begins. . .”: title chosen by one of the TAs.
- “Speaking in Tongues”: title of an album by the Talking Heads, and a song by Arcade Fire.

2 The trouble with substitution

We’ve defined dynamic semantics in two different ways: big-step and small-step. In both, we used substitution to define what expressions with identifiers (variables) mean: a with expression evaluates its bound expression, and immediately replaces all the instances of the bound identifier with that expression. Functions (lam/app) and recursion (rec) also were given meaning via substitution.

Our notion of substitution is directly descended from Church’s λ -calculus, but as a general (and less precise) notion, substitution is older: in algebra we can substitute 5 for x in

$$x^2 + 3$$

to get $5^2 + 3$. (I don’t think the ancient syllogisms of Greece and India—“Socrates is a man, all men are mortal, therefore Socrates is mortal”; “This hill is smoky; whatever is smoky is fiery (for example: a kitchen); therefore this hill is fiery”¹—are truly *substitution*: there are no variables.)

The connection to the λ -calculus, which was shown to be equivalent in power to Turing machines, guarantees that substitution is a “right way” of defining how features like with and app work. It does not mean that substitution is *the* right way of defining how those features work. In fact, substitution is (almost?) never used to implement interpreters.

Substitution has several disadvantages, compared to other methods:

- *Inefficiency*: Every time our interpreter calls $subst(e1, x, e2)$, our implementation of $subst$ searches for $(id\ x)$ throughout the entire expression $e1$. It must do this even if $e1$ is very large, and $(id\ x)$ appears just once (or even not at all!).
- *Obscurity*: Giving a function (or other expression) a name is important for clarity and convenience; we would rather write `{app double 5}` than `{app {lam y {+ y y} 5}}`, even though they give the same result. But a substitution-based interpreter that prints the expressions it’s evaluating (like `visible-interp.rkt` does) will show you the latter. This is perhaps most aggravating with recursive functions.

Against these, we should weigh substitution’s advantages:

- *Simplicity*: The definition of substitution is more concise and straightforward than other methods.
- *Versatility*: While substitution doesn’t “scale” in terms of performance (see “Inefficiency” above), it “scales up” well across a variety of language features. The same, relatively simple, style of defining substitution works for languages with functions that return functions (“first-class” functions) and for recursive functions. Environments are more brittle: adding new features sometimes requires us to define environment in a way that is more complicated (rather than just being *longer*, as is the case with substitution).

Whether or not you prefer environments, you should learn about them, especially if you plan to take CPSC 411.

¹Adapted from Vidyabhusana, *A History of Indian Logic* (1920), p. 61.

3 Environments

The idea of environment-based dynamic semantics is that, to evaluate (with $x \in e_1$), we won't evaluate e_2 to v_2 and then substitute v_2 for x ; instead, we will evaluate e_2 to v_2 , and “remember” the fact that x has the value v_2 . This fact will be stored in an *environment* that maps identifiers to values. If and when we need to evaluate an instance of x in the body e_1 , that is, if we need to evaluate $(\text{id } x)$, we look up x and use the value we find, which will be v_2 .

In a sense, we are simulating substitution: if we had substituted v_2 for x , we would find v_2 inside the body e_1 .

3.1 Back to basics: WAE

Because environments are more brittle than substitution, I think it's better to “roll back” our Fun language to WAE (arithmetic expressions and with), define the simplest possible environments, and then carefully evolve our notion of an environment as we restore language features.

Quoting 04-operational.pdf:

$$\begin{aligned} \langle \text{WAE} \rangle ::= & \langle \text{num} \rangle \\ & | \{ + \langle \text{WAE} \rangle \langle \text{WAE} \rangle \} \\ & | \{ - \langle \text{WAE} \rangle \langle \text{WAE} \rangle \} \\ & | \{ \text{with } \{ \langle \text{id} \rangle \langle \text{WAE} \rangle \} \langle \text{WAE} \rangle \} \\ & | \langle \text{id} \rangle \end{aligned}$$

```
(define-type WAE
  [num (n number?)]
  [add (lhs WAE?) (rhs WAE?)]
  [sub (lhs WAE?) (rhs WAE?)]
  [with (name symbol?) (named-expr WAE?) (body WAE?)]
  [id (name symbol?)])
```

§3 Environments

3.2 Mapping identifiers to expressions

To get an idea of what is needed, consider the WAE expression (in abstract syntax)

$$(\text{with } x \text{ (num 3) (with } y \text{ (num 4) (add (id } x \text{) (id } y \text{))))$$

If we don't use substitution, when we evaluate $(\text{add (id } x \text{) (id } y \text{)})$ we need to remember that x was bound to (num 3) , and y was bound to (num 4) . We need a “lookup table” that maps identifiers to expressions.

In Typed Fun, we used a typing context Γ to map identifiers to types, and defined what those contexts were with a grammar:

$$\begin{array}{ll} \text{Typing contexts } \Gamma ::= \emptyset & \text{empty context (no assumptions)} \\ | x : A, \Gamma & x \text{ has type } A, \text{ with more assumptions} \end{array}$$

We'll do the same for environments:

$$\begin{array}{ll} \text{Environments (for WAE) } env ::= \emptyset & \text{empty environment} \\ | x=e, env & x \text{ bound to } e, \text{ with “more environment”} \end{array}$$

(It would be more standard to use the Greek letter rho (ρ), rather than “env”, but we've used enough Greek letters for now.)

For consistency with typing contexts Γ , environments env will grow to the left, like cons-lists in Racket.

Let's consider an even smaller example than the one above.

$$(\text{with } y \text{ (num 4) (add (num 3) (id } y \text{)))}$$

If this expression is the entire program, it's not inside any withs, so the environment env is empty when we start evaluating.

Regardless of how environments work, (num 4) should still evaluate to (num 4) . But now we need to remember that y is (num 4) , so we'll need to evaluate the body $(\text{add (num 3) (id } y \text{)})$ under the environment

$$y=(\text{num 4}), \emptyset$$

(It's okay to write this as just $y=(\text{num 4})$; here, I want to emphasize that we started with \emptyset , and are growing the environment leftwards.)

Then, while evaluating $(\text{id } y \text{)}$ in $(\text{add (num 3) (id } y \text{)})$, we will look up $(\text{id } y \text{)}$ in the current environment $y=(\text{num 4}), \emptyset$, and evaluation will behave as if we were evaluating $(\text{add (num 3) (num 4)})$.

Just as we used Γ in the typing judgment $\Gamma \vdash e : A$, we'll use env in a new *environment-based evaluation* judgment form

$$env \vdash e \Downarrow v$$

We'll also assume that a “lookup function” $lookup(env, x)$ has been defined, so that

$$lookup(env, x) = e$$

if the environment env contains $x=e$. (In our Racket code, we have a function `look-up-id`.)

§3 Environments

I think we have enough to revise the evaluation rules for WAE. What were those?

$e \Downarrow v$ WAE expression e evaluates to value v

$$\frac{}{(\text{num } n) \Downarrow (\text{num } n)} \text{Eval-num}$$

$$\frac{e1 \Downarrow (\text{num } n1) \quad e2 \Downarrow (\text{num } n2)}{(\text{add } e1 \ e2) \Downarrow (\text{num } n1 + n2)} \text{Eval-add}$$

$$\frac{e1 \Downarrow (\text{num } n1) \quad e2 \Downarrow (\text{num } n2)}{(\text{sub } e1 \ e2) \Downarrow (\text{num } n1 - n2)} \text{Eval-sub}$$

$$\frac{e1 \Downarrow v1 \quad \text{subst}(e2, x, v1) \Downarrow v2}{(\text{with } x \ e1 \ e2) \Downarrow v2} \text{Eval-with}$$

$$\frac{}{(\text{id } x) \text{ free-variable-error}} \text{Eval-free-identifier}$$

$\text{env} \vdash e \Downarrow v$ Under environment env ,
WAE expression e evaluates to value v

$$\frac{}{\text{env} \vdash (\text{num } n) \Downarrow (\text{num } n)} \text{Env-num}$$

$$\frac{e1 \Downarrow (\text{num } n1) \quad e2 \Downarrow (\text{num } n2)}{(\text{add } e1 \ e2) \Downarrow (\text{num } n1 + n2)} \text{Env-add}$$

$$\frac{e1 \Downarrow (\text{num } n1) \quad e2 \Downarrow (\text{num } n2)}{(\text{sub } e1 \ e2) \Downarrow (\text{num } n1 - n2)} \text{Env-sub}$$

$$\frac{e1 \Downarrow v1 \quad \Downarrow v2}{(\text{with } x \ e1 \ e2) \Downarrow v2} \text{Env-with}$$

$$\frac{}{(\text{id } x) \Downarrow} \text{Env-id}$$

$$\frac{}{\text{unknown-id-error}} \text{Env-unknown-id}$$

■ **Exercise 1.** (Do it tonight, before class, if feasible.)

I left some blank space in the “Env-...” rules. Fill it in with whatever is needed. Env-num is finished, and you can follow that pattern for some of the other rules.

If you’re not sure how to start, I already updated *part* of the function `env-interp` in `env-with-broken.rkt` (link on the notes page) to reflect how I would fill in Env-add and Env-sub, so you can map back from that code if you like. But I haven’t written the code for the more interesting rules yet...

The completed rules are on the next page.

§3 Environments

$e \Downarrow v$ WAE expression e evaluates to value v

$$\frac{}{(\text{num } n) \Downarrow (\text{num } n)} \text{Eval-num}$$

$$\frac{e1 \Downarrow (\text{num } n1) \quad e2 \Downarrow (\text{num } n2)}{(\text{add } e1 \ e2) \Downarrow (\text{num } n1 + n2)} \text{Eval-add}$$

$$\frac{e1 \Downarrow (\text{num } n1) \quad e2 \Downarrow (\text{num } n2)}{(\text{sub } e1 \ e2) \Downarrow (\text{num } n1 - n2)} \text{Eval-sub}$$

$$\frac{e1 \Downarrow v1 \quad \text{subst}(e2, x, v1) \Downarrow v2}{(\text{with } x \ e1 \ e2) \Downarrow v2} \text{Eval-with}$$

$$\frac{}{(\text{id } x) \text{ free-variable-error}} \text{Eval-free-identifier}$$

$\text{env} \vdash e \Downarrow v$ Under environment env , WAE expression e evaluates to value v

$$\frac{}{\text{env} \vdash (\text{num } n) \Downarrow (\text{num } n)} \text{Env-num}$$

$$\frac{\text{env} \vdash e1 \Downarrow (\text{num } n1) \quad \text{env} \vdash e2 \Downarrow (\text{num } n2)}{\text{env} \vdash (\text{add } e1 \ e2) \Downarrow (\text{num } n1 + n2)} \text{Env-add}$$

$$\frac{\text{env} \vdash e1 \Downarrow (\text{num } n1) \quad \text{env} \vdash e2 \Downarrow (\text{num } n2)}{\text{env} \vdash (\text{sub } e1 \ e2) \Downarrow (\text{num } n1 - n2)} \text{Env-sub}$$

$$\frac{\text{env} \vdash e1 \Downarrow v1 \quad x=v1, \text{env} \vdash e2 \Downarrow v2}{\text{env} \vdash (\text{with } x \ e1 \ e2) \Downarrow v2} \text{Env-with}$$

$$\frac{\text{lookup}(\text{env}, x) = e}{\text{env} \vdash (\text{id } x) \Downarrow e} \text{Env-id}$$

$$\frac{\text{lookup}(\text{env}, x) \text{ undefined}}{\text{env} \vdash (\text{id } x) \text{ unknown-id-error}} \text{Env-unknown-id}$$

3.3 The Shadow Chancellor² Strikes Back

With substitution, we saw that expressions that repeatedly bind the same identifier are evaluated with the inner binding “shadowing” the outer one, so that

$$(\text{with } x \ (\text{num } 1) \ (\text{with } x \ (\text{num } 2) \ (\text{id } x)))$$

evaluates to $(\text{num } 2)$, not $(\text{num } 1)$. The environment-based semantics will behave the same way, but only because of a particular way we’re defining *lookup*: it starts looking from the left.

$$\frac{\frac{\frac{}{\emptyset \vdash (\text{num } 1)}{\Downarrow (\text{num } 1)} \text{Env-num} \quad \frac{\frac{}{x=(\text{num } 1), \emptyset \vdash (\text{num } 2)}{\Downarrow (\text{num } 2)} \text{Env-num}}{x=(\text{num } 1), \emptyset \vdash (\text{with } x \ (\text{num } 2) \ (\text{id } x)) \Downarrow (\text{num } 2)} \text{Env-with}}{\emptyset \vdash (\text{with } x \ (\text{num } 1) \ (\text{with } x \ (\text{num } 2) \ (\text{id } x))) \Downarrow (\text{num } 2)} \text{Env-with}}{\text{lookup}((x=(\text{num } 2), x=(\text{num } 1), \emptyset), x) = (\text{num } 2)) \text{Env-id}} \quad \frac{\text{lookup}((x=(\text{num } 2), x=(\text{num } 1), \emptyset), x) = (\text{num } 2))}{x=(\text{num } 2), x=(\text{num } 1), \emptyset \vdash (\text{id } x) \Downarrow (\text{num } 2)} \text{Env-id}}{\text{Env-with}} \text{Env-with}$$

We should really define *lookup* using rules; I’ll leave that as an exercise (next page).

■ **Exercise 2.** Fill in the rules below, which derive a judgment $\text{lookup}(\text{env}, x) = e$: (Feel free to translate “backwards” from the Racket implementation of `look-up-id`.)

²At some point, the UK Parliament becomes indistinguishable from a bad fantasy novel.

$$\frac{}{lookup(\emptyset, x) = \dots} \text{lookup-empty}$$

$$\frac{}{lookup((x=e, env), x) = \dots} \text{lookup-found} \quad \frac{}{lookup((y=e, env), x) = \dots} \text{lookup-next}$$

3.4 Question Period

Question:

The expression after the “ \Downarrow ” should always be a value. Shouldn’t Env-id evaluate e to v ?

It could, but it doesn’t need to: the expressions we put into environments are all values. The only rule that adds anything to the environment is Env-with, and the expression it adds is $v1$, which is a value.

Question: Could Env-with *not* evaluate $e1$, and put $e1$ into the environment, instead?

In that case, Env-id *would* need to evaluate the expression it gets from *lookup*. That would give us an “expression strategy” for with. That’s inconsistent with our substitution-based semantics, but it’s not wrong; it’s just not what I want to do.

Question: What if Env-with puts $e1$ into the environment, and Env-id evaluates that expression to get $v1$, and then *updates the environment* with $v1$? Would that give us lazy evaluation?

You could certainly implement that—for example, using Racket’s mutable “boxes”. Moreover, we could model it using rules. But the rules would need to be rather different from the above rules, which derive the judgment form $env \vdash e \Downarrow v$. That judgment can’t model a change to the environment; the above rules can only add to the environment *inside* a premise. So if your environment is

$$\underbrace{x=(\text{add } (\text{num } 1) (\text{num } 1)), \emptyset}_{env}$$

and you evaluate $(\text{add } (\text{id } x) (\text{id } x))$, you can’t “transmit” the updated *env* from the first premise to the second premise. Rules and derivations aren’t mutable.

$$\frac{env \vdash (\text{id } x) \Downarrow (\text{num } 2) \quad env \vdash (\text{id } x) \Downarrow (\text{num } 2)}{\underbrace{x=(\text{add } (\text{num } 1) (\text{num } 1)), \emptyset}_{env} \vdash (\text{add } (\text{id } x) (\text{id } x)) \Downarrow (\text{num } 4)} \text{Env-add}$$

However, you could change the judgment form to something like

$$env \vdash e \Downarrow v, env'$$

which could be read “starting in environment env , evaluating expression e produces value v and an environment env' .” Then the conclusion of the rule for *id* could have the “updated” environment as env' .

We’ll need to do something like this to model mutable state (hopefully, next week).