

Rules of Evaluation of OCaml Expressions

P. Madhusudan

September 2018

1 A fragment of Ocaml

Expressions: Let us define a fragment of Ocaml for which we can write formal rules of evaluation.

$$e ::= c \mid x \mid (e_1, \dots, e_n) \mid op\ e \mid e_1\ e_2 \mid fun\ x \rightarrow e \mid let\ (x_1, \dots, x_n) = (e_1, \dots, e_n)\ in\ e \\ \mid if\ e\ then\ e_1\ else\ e_2 \\ \mid match\ e_0\ with\ p_1 \rightarrow e_1 \mid \dots \mid p_n \rightarrow e_n$$

In the above, e, e_i denote an expressions. c denotes constants, i.e., concrete values in various datatypes (like 23, 3.1415, "bilbo", etc.). op is a built-in primitive operator (like + or +. or ^), interpreted as being curried and hence taking only one parameter at a time. x and f are identifiers, and p_i are patterns.

Note that 1-tuples (e_1) are synonymous with e_1 . Hence "let $x_1 = e_1$ in e " is written really as "let $(x_1) = (e_1)$ in e ", and we avoid a special syntax for it.

Note that the above fragment disallows recursion ("let rec"), algebraic datatypes, etc. It also does not allow pattern matching in "let" statements (this can be accommodated easily, though).

2 Evaluation

Values Let us define the set of *values* (*Values*) as a subset of expressions (as defined above) and *closures*:

$$v, v_i ::= c \mid (v_1, \dots, v_n) \mid \langle x \rightarrow e, \rho \rangle$$

Environment Let us define environments as partial functions $\rho : Ids \rightarrow Values$. Environments have finite domains.

Let us also define *updates* on environments. Given two environments ρ_1 and ρ_2 , with domains D_1 and D_2 respectively, the domain of $\rho_1 + \rho_2$ is $D_1 \cup D_2$, and for each literal $x \in D_1 \cup D_2$,

$$(\rho_1 + \rho_2)(x) = \rho_1(x)\ if\ x \in D_1 \\ = \rho_2(x)\ otherwise\ i.e.,\ if\ x \in D_2 \setminus D_1$$

The above says that the environment ρ_2 is updated by ρ_1 's definitions, and the definitions given in ρ_1 supersede those in ρ_2 .

Evaluation We can now define the evaluation function $Eval$ that maps pairs, each pair consisting of an expression and an environment, to a value. The notion of an evaluation, strictly speaking, will be a partial map in general (evaluation will not be defined if for example an expression uses a variable that is not defined in the environment, or if when we have recursive functions, the evaluation of the functions do not terminate).

We will also define a function App that evaluates a closure on an expression to a value. These two functions, $Eval$ and App are mutually recursively defined.

These recursive definitions will be defined using *equations*. Unlike many recursive definitions you typically write, there is no *measure* that decreases (i.e., when defining the evaluation of an expression, the definition may involve a recursive definition that could be on *larger* expressions). This does not happen in the definitions below, as we do not handle evaluation of recursive functions, but they will happen in general. The semantics of these equations is hence defined using what is called *least-fixed point semantics*— we mean that the *least* partial map that satisfies the equations is the meaning of $Eval$ and App . You may wonder why such a least map must exist— it can be shown to exist using Tarski-Knaster theorem, which says the least fixpoint exists because the right-hand-side of these definitions are *monotonic*. Let’s not worry about these for now.

$$\begin{aligned}
Eval(c, \rho) &= c \\
Eval(x, \rho) &= \rho(x), \quad \text{if } \rho(x) \text{ is defined} \\
&= \text{undefined, otherwise} \\
Eval((e_1, \dots, e_n), \rho) &= (v_1, \dots, v_n), \quad \text{where } Eval(e_n, \rho) = v_n, \dots, Eval(e_1, \rho) = v_1 \\
Eval(op \ e) &= [[op]](v), \quad \text{where } Eval(e) = v \\
Eval(e_1 \ e_2, \rho) &= App(Eval(e_1, \rho), Eval(e_2, \rho)) \\
Eval(fun \ x \rightarrow e, \rho) &= \langle x \rightarrow e, \rho \rangle \\
Eval(let \ (x_1, \dots, x_n) = (e_1, \dots, e_n) \ in \ e, \rho) &= Eval(e, \{x_1 \rightarrow v_1, \dots, x_n \rightarrow v_n\} + \rho) \\
&\quad \text{where for each } i, Eval(e_i, \rho) = v_i \\
Eval(if \ e \ then \ e_1 \ else \ e_2, \rho) &= Eval(e_1, \rho) \quad \text{if } Eval(e, \rho) = \text{true} \\
&= Eval(e_2, \rho) \quad \text{if } Eval(e, \rho) = \text{false} \\
Eval(match \ e \ with \ p_1 \rightarrow e_1 \mid \dots \mid p_n \rightarrow e_n, \rho) &= Eval(e_1, b_1 + \rho) \quad \text{if } Eval(e, \rho) \text{ matches } p_1 \text{ producing binding } b_1 \\
&= \dots \\
&= Eval(e_n, b_n + \rho) \quad \text{if } Eval(e, \rho) \text{ matches } p_n \text{ producing binding } b_n \\
&= \text{undefined} \quad \text{otherwise} \\
App(\langle (x_1, \dots, x_m) \rightarrow e, \rho \rangle, (v_1, \dots, v_m)) &= Eval(e, \{x_1 \rightarrow v_1, \dots, x_m \rightarrow v_m\} + \rho)
\end{aligned}$$

Some remarks on the above definitions, line by line:

- Constants evaluate to themselves.
- A variable evaluates to the value given to it by the environment.
- A tuple of expressions evaluates to a tuple of values obtained by evaluating the individual expressions. Note that we don’t specify above the order in which these expressions are evaluated— they are typically evaluated right-to-left, or the order may be unspecified, in which case we assume compilers can evaluate them in any order and the programmer cannot assume a particular order.
- A primitive operation op is evaluated by calling a pre-existing implementation of the function associated with op , which we call $[[op]]$.

- When evaluating a function application $(e_1 e_2)$, we first evaluate e_1 (to get a closure) and evaluate e_2 , using the current environment, and then call *App* on the resulting values. Note that *App* itself does not get the environment ρ (it will have access to the environment in the closure for evaluating the body of the function, but does not need the current environment ρ).
- Evaluation of an anonymous function definition evaluates to its closure, where the closure imbibes the current environment.
- In evaluating a *let* $(x_1, \dots, x_n) = (e_1, \dots, e_n)$ in e expression, we first evaluate the expressions e_1, \dots, e_n (in some order, but Ocaml usually does this in reverse order) and then evaluate the expression e in the environment that has been updated with the bindings of each x_i to the corresponding value.
- In evaluating a conditional, the condition e is first evaluated. Assuming this evaluation is well defined (terminates), we evaluate e_1 if the condition evaluates to true and evaluate e_2 if the condition evaluates to false. Note that if the condition evaluates to true, e_2 is *not* evaluated and, similarly, e_1 is not evaluated if e evaluates to false.
- In evaluating a match statement, we evaluate e_0 first, and then find the first pattern p_i that matches and find the binding that it results in. We then evaluate (only) the expression e_i corresponding to the pattern p_i , with the environment updated with the binding b_i . We don't describe here formally how matching is done.
- The *App* function takes a closure and a tuple of values and applies the function body in the closure by updating the environment in the closure (which is the environment at the time of creation of the function) with the given values for the parameter to the function.