

Lambda-calculus and programming language semantics

Thibaut Balabonski @ UPSay

Winter 2023

<https://public.lmf.cnrs.fr/~blsk/LambdaCalculus/>

Chapter 5: implementing the λ -calculus

1 Interpretation

Following the operational semantics of the λ -calculus, we can define an *interpreter*, that is a function that takes as input a λ -term and returns the value computed by this term.

Terms:

$t ::= c$	constant
x	variable
$\lambda x.t$	abstraction
$t u$	application

Values:

$v ::= c$
$\lambda x.t$

Call by value, small-step semantics. We define a reduction relation $t \rightarrow t'$ describing one step of reduction. Call-by-value evaluation: evaluate the argument before resolving a function application. The first rule only applies if the argument is a value v .

$$\frac{}{(\lambda x.t) v \rightarrow t\{x \leftarrow v\}} \quad \frac{t \rightarrow_{\beta} t'}{t u \rightarrow t' u} \quad \frac{u \rightarrow_{\beta} u'}{v u \rightarrow v u'}$$

The last rule only applies if the left member of the application is a value v . This forces evaluation from left to right. Note that we never reduce under λ -abstractions, and that we do not expect to reach a free variable.

Evaluation is done through a sequence of such steps:

$$t \rightarrow^* v$$

We could define an interpreter performing sequences of steps, following this definition. The cost of one step would contain:

- the cost of finding the next redex $(\lambda x.t) u$ (proportional to the depth of the redex),
- the cost of performing the substitution (proportional to the size of t) and managing required renamings,
- the cost of reconstructing the term after transforming the redex (possibly proportional to the depth of the redex).

This cost would then be repeated for each step. This is *very* inefficient.

Call-by-value, big-step semantics. We define an evaluation relation $t \Downarrow v$ defining the value v resulting from the evaluation of t . The relation is reflexive on terms that are already values.

$$\frac{}{c \Downarrow c} \quad \frac{}{\lambda x.t \Downarrow \lambda x.t} \quad \frac{t \Downarrow \lambda x.s \quad u \Downarrow v \quad s\{x \leftarrow v\} \Downarrow w}{t u \Downarrow w}$$

As before, we never reduce under λ -abstractions, and we do not expect to reach a free variable. Evaluation of an application requires evaluating the function and the argument, and then evaluating the function body, with access to the value of the argument.

Following these rules, we define an interpretation function eval .

$$\begin{aligned} \text{eval}(c) &= c \\ \text{eval}(\lambda x.t) &= \lambda x.t \\ \text{eval}(t u) &= \text{eval}(s\{x \leftarrow v\}) \quad \text{if } \text{eval}(t) = \lambda x.s \text{ and } \text{eval}(u) = v \end{aligned}$$

The costs of finding the next redex and of rebuilding the term are amortized. All the costs related to substitution are still there.

Call-by-value, big-step semantics with environments. Rather than actually performing substitutions, we consider a term (with free variables) together with an environment (which gives values to the free variables of the term). The evaluation relation now takes the form

$$e \vdash t \Downarrow v$$

and means “the term t taken in the environment e evaluates to v ”.

To correctly implement the scoping rules of variables in the λ -calculus, every value possibly containing free variables should be given together with the environment in which it has been defined. Thus, the “value” associated to a λ -abstraction $\lambda x.t$ is a *closure*, that is a pair $\langle \lambda x.t, e \rangle$ of the λ -abstraction and an environment e providing values for its free variables.

The values and environments are now defined by:

$$\begin{aligned} v &::= c \\ &\quad | \langle \lambda x.t, e \rangle \\ e &::= \{x_1 \mapsto v_1, \dots, x_k \mapsto v_k\} \end{aligned}$$

Evaluation rules are as follows. The value of a variable is looked up in the environment. The value of a λ -abstraction is a closure built with the current environment. When evaluating the body of a function, we use the environment provided by the closure, to ensure that all free variables are associated to the proper (lexically scoped) values. We write $\{e, x \mapsto v\}$ the extension of the environment e with the mapping of the variable x to the value v .

$$\begin{array}{c} \frac{}{e \vdash c \Downarrow c} \qquad \frac{e(x) = v}{e \vdash x \Downarrow v} \qquad \frac{}{e \vdash \lambda x.t \Downarrow \langle \lambda x.t, e \rangle} \\ \frac{e \vdash t \Downarrow \langle \lambda x.s, e' \rangle \quad e \vdash u \Downarrow v \quad \{e', x \mapsto v\} \vdash s \Downarrow w}{e \vdash t u \Downarrow w} \end{array}$$


Following these rules, we define a new interpretation function, that takes as inputs a λ -term and an environment.

$$\begin{aligned} \text{eval}(c, e) &= c \\ \text{eval}(x, e) &= e(x) \\ \text{eval}(\lambda x.t, e) &= \langle \lambda x.t, e \rangle \\ \text{eval}(t u, e) &= \text{eval}(s, \{e', x \mapsto v\}) \quad \text{if } \text{eval}(t, e) = \langle \lambda x.s, e' \rangle \text{ and } \text{eval}(u, e) = v \end{aligned}$$

There is no substitution anymore. However we still have to deal with names.

Call-by-value, big-step semantics with environments, without names. We can represent λ -terms without using variables names, by replacing each variable occurrence by a *number* called *de Bruijn*

index, defined as the number of λ that are between this variable occurrence and its binder.



With this nameless notations for terms, a β -redex now has the shape

$$(\lambda.t) \ u$$

and its reduction consists in replacing by u the variable occurrences of t that are bound by this λ . They are precisely: the occurrences of 0 that are not nested under any other λ , the occurrences of 1 that are nested under exactly 1 other λ , the occurrences of 2 that are nested under exactly 2 other λ 's... However, in our implementation we do not perform any substitution: the argument u is just inserted at index 0 in the environment. More generally, an environment is now just a list of values, and each value is accessed through its index in the list.

Terms, values and environments. We write \underline{n} for the variable with index n . An environment is now written as a sequence of values (there are no names anymore, we refer to values in the environment only by their position).

$t ::= c$	constant
$\quad \ \underline{n}$	variable (de Bruijn index)
$\quad \ \lambda.t$	abstraction
$\quad \ t \ u$	application
$v ::= c$	
$\quad \ \langle \lambda.t, e \rangle$	
$e ::= v_0 \cdot v_1 \cdot \dots \cdot v_k$	

Evaluation rules. We write $e[n]$ for taking the value of index n in the list e . We write $v \cdot e$ for the environment obtained from by inserting v at index 0 (and incrementing the indices of all the other values in e).

$$\frac{}{e \vdash c \Downarrow c} \quad \frac{e[n] = v}{e \vdash \underline{n} \Downarrow v} \quad \frac{}{e \vdash \lambda.t \Downarrow \langle \lambda.t, e \rangle} \quad \frac{e \vdash t \Downarrow \langle \lambda.s, e' \rangle \quad e \vdash u \Downarrow v \quad v \cdot e' \vdash s \Downarrow w}{e \vdash t \ u \Downarrow w}$$

These rules give us a final version of our evaluation function, working on λ -terms in nameless representation.

$$\begin{aligned} \text{eval}(c, e) &= c \\ \text{eval}(\underline{n}, e) &= e[n] \\ \text{eval}(\lambda.t, e) &= \langle \lambda.t, e \rangle \\ \text{eval}(t \ u, e) &= \text{eval}(s, v \cdot e') \quad \text{if } \text{eval}(t) = \langle \lambda.s, e' \rangle \text{ and } \text{eval}(u, e) = v \end{aligned}$$

Efficient interpreter. The final evaluation function can easily be translated into code, giving us an efficient interpreter (an interpreter from which the obvious inefficiencies have been removed). In the following caml code we assume a datatype 'a env for environments containing elements of type 'a, with functions env_lookup: int -> 'a env -> 'a for getting the element at a given index, and env_add: 'a -> 'a env -> 'a env for inserting an element at index 0. Many concrete data structures exist, which provide efficient implementations of these operations.

```

type term =
  | Cst of int
  | Var of int
  | Abs of term
  | App of term * term
type value =
  | VCst of int
  | VClos of term * value env

let rec eval e env = match e with
  | Cst n -> VCst n
  | Var n -> env_lookup n env
  | Abs t -> VClos(t, env)
  | App(t, u) ->
    let VClos(s, env') = eval t env in
    let v = eval u env in
    eval s (env_add v env')

```

2 Compilation to abstract machine code

Instead of interpreting a program, we can translate it into a sequence of instructions for some machine. Two possibilities:

- we can target an actual microprocessor, and the instructions will be executed directly by the hardware,
- or we can target an *abstract machine*, which looks like a machine but does not correspond to actual hardware.

In general, the instructions of an abstract machine are chosen with two criteria: they are close to the basic mechanisms of the source language, and they can be implemented efficiently.

Warm-up: an abstract machine for evaluating arithmetic expressions. Arithmetic expression can famously be evaluated using a very simple machine working with a *stack* (a sequential data-structure in which we can add and remove elements “at the top”).

Arithmetic expressions:

$$\begin{array}{l}
 a ::= N \\
 \quad | \quad a_1 + a_2 \\
 \quad | \quad a_1 \times a_2
 \end{array}$$

Instructions:

- $\text{CONST}(N)$: push integer N on top of the stack,
- ADD , MUL : pop two integers from the stack, combine, push the result back on the stack,
- $c_1; c_2$: execute c_1 , then c_2 .

An expression a is translated into a sequence of instructions $C(a)$ that evaluates a and pushes the obtained value on the stack. Evaluating a binary operation consists in evaluating both parts, saving the intermediate values on the stack, and then combining the two stored values.

$$\begin{array}{l}
 C(N) = \text{CONST}(N) \\
 C(a_1 + a_2) = C(a_1); C(a_2); \text{ADD} \\
 C(a_1 \times a_2) = C(a_1); C(a_2); \text{MUL}
 \end{array}$$

The state of the machine executing such instructions is defined by:

- a code pointer, to the instructions that have to be executed,
- a stack with the stored intermediate results.

The execution of the machine is a sequence of steps, defined by the following table.

State before		State after	
code	stack	code	stack
CONST(N); c	s	c	$N \cdot s$
ADD; c	$N_1 \cdot N_2 \cdot s$	c	$(N_1 + N_2) \cdot s$

This machine can be implemented efficiently, either by interpreting its instructions, or by expanding them into native code for our actual processor.

An abstract machine for call-by-value λ -calculus. The “SECD” machine for executing λ -terms has three elements:

- a code pointer c , to the instructions that have to be executed,
- an environment e , which gives values to the variables,
- a stack s , which stores intermediate results and other useful context information.

These components are the C, E and S from SECD. The original version contained a second stack (“dump”), merged here into S. In addition to the basic arithmetic operation, it features the following instructions:

- ACCESS(n): push the n -th element of the environment on the stack,
- CLOSURE(c): make a closure with code c and the current environment, and push it on the stack,
- APPLY: pop an argument and a function closure, and perform the application,
- RETURN: returns to caller.

Compilation:

$$\begin{aligned}
 C(n) &= \text{ACCESS}(n) \\
 C(\lambda.t) &= \text{CLOSURE}(C(t); \text{RETURN}) \\
 C(t \ u) &= C(t); C(u); \text{APPLY}
 \end{aligned}$$

Execution steps. The APPLY instruction records on the stack the context information that should be restored after the call. This information consists in the code pointer c to the instructions that follow the code, and the current environment. Then the RETURN instruction retrieves and restores this information. We write again $\langle c, e \rangle$ for a closure with code c and environment e (note that here, c is a sequence of instructions and not a λ -term). Notation (c, e) similarly denotes a pair of code an environment, used to restore context after function return.

State before			State after		
code	env	stack	code	env	stack
ACCESS(n); c	e	s	c	e	$e[n] \cdot s$
CLOSURE(c'); c	e	s	c	e	$\langle c', e \rangle \cdot s$
APPLY; c	e	$v \cdot \langle c', e' \rangle \cdot s$	c'	$v \cdot e'$	$(c, e) \cdot s$
RETURN; c	e	$v \cdot (c', e') \cdot s$	c'	e'	$v \cdot s$

Compilation, extended for (nameless) local variables. Since it is used a lot, the definition of local variables using let is given a pair of dedicated instructions. The scope of the local variable is defined by a pair LET/ENDLET that adds the value to the environment, and removes it at the end.

$$C(\text{let } t \text{ in } u) = C(t); \text{LET}; C(u); \text{ENDLET}$$

- LET: add to the environment the value taken at the top of the stack,
- ENDLET: drop a value from the environment.

Execution steps.

State before			State after		
code	env	stack	code	env	stack
LET; c	e	$v \cdot s$	c	$v \cdot e$	s
ENDLET; c	$v \cdot e$	s	c	e	s

Proving partial correctness of the machine. Goal: if a term t evaluates to v in the λ -calculus (meaning: if $t \rightarrow^* v$), then the machine started in the state $\langle C(t) \mid \varepsilon \mid \varepsilon \rangle$ terminates in a state $\langle \varepsilon \mid \varepsilon \mid v' \rangle$ where v' represents the value v .

We prove this by induction on the big-step semantics $e \vdash t \Downarrow v$. For the induction to work, we need to generalize our goal property in such a way that it will be applicable to the evaluation of subterms. For this, we include in the statement an arbitrary sequence of code, an environment and a stack. The property we prove is then:

$$\text{If } e \vdash t \Downarrow v \text{ then for any } c \text{ and } s \text{ we have } \langle C(t); c \mid \mathcal{T}(e) \mid s \rangle \rightarrow^* \langle c \mid \mathcal{T}(e) \mid \mathcal{T}(v) \cdot s \rangle$$

This property uses a translation function $\mathcal{T}(\cdot)$ for translating values and environments from the λ -calculus in machine format.

$$\begin{aligned} \mathcal{T}(N) &= N \\ \mathcal{T}(\langle \lambda.t, e \rangle) &= \langle C(t); \text{RETURN}, \mathcal{T}(e) \rangle \\ \mathcal{T}(v_0 \cdot v_1 \cdot \dots \cdot v_k) &= \mathcal{T}(v_0) \cdot \mathcal{T}(v_1) \cdot \dots \cdot \mathcal{T}(v_k) \end{aligned}$$

Proof of the property, by induction over the derivation of $e \vdash t \Downarrow v$.

- Case $e \vdash N \Downarrow N$. Then $C(N) = \text{CONST}(N)$ and $\mathcal{T}(N) = N$ and we have

$$\langle \text{CONST}(N); c \mid \mathcal{T}(e) \mid s \rangle \rightarrow \langle c \mid \mathcal{T}(e) \mid N \cdot s \rangle$$

in one step of the machine.

- Case $e \vdash \underline{n} \Downarrow v$, with $e[\underline{n}] = v$. Then $C(\underline{n}) = \text{ACCESS}(n)$ and $\mathcal{T}(e)[\underline{n}] = \mathcal{T}(v)$ and we have

$$\langle \text{ACCESS}(N); c \mid \mathcal{T}(e) \mid s \rangle \rightarrow \langle c \mid \mathcal{T}(e) \mid \mathcal{T}(v) \cdot s \rangle$$

in one step of the machine.

- Case $e \vdash \lambda.t \Downarrow \langle \lambda.t, e \rangle$. Then $C(\lambda.t) = \text{CLOSURE}(C(t); \text{RETURN})$ and $\mathcal{T}(\langle \lambda.t, e \rangle) = \langle C(t); \text{RETURN}, \mathcal{T}(e) \rangle$ and we have

$$\langle \text{CLOSURE}(C(t); \text{RETURN}); c \mid \mathcal{T}(e) \mid s \rangle \rightarrow \langle c \mid \mathcal{T}(e) \mid \langle C(t); \text{RETURN}, \mathcal{T}(e) \rangle \cdot s \rangle$$

in one step of the machine.

- Case $e \vdash t u \Downarrow w$ with $e \vdash t \Downarrow \langle \lambda.r, e' \rangle$, $e \vdash u \Downarrow v$ and $v \cdot e' \vdash r \Downarrow w$. Then $C(t u) = C(t); C(u); \text{APPLY}$ and $\mathcal{T}(\langle \lambda.r, e' \rangle) = \langle C(r); \text{RETURN}, \mathcal{T}(e') \rangle$. We build the following sequence of execution of the machine:

$$\begin{array}{l} \rightarrow^* \\ \rightarrow^* \\ \rightarrow \\ \rightarrow^* \\ \rightarrow \end{array} \left| \begin{array}{l} C(t); C(u); \text{APPLY}; c \\ C(u); \text{APPLY}; c \\ \text{APPLY}; c \\ C(r); \text{RETURN} \\ \text{RETURN} \\ c \end{array} \right| \left| \begin{array}{l} \mathcal{T}(e) \\ \mathcal{T}(e) \\ \mathcal{T}(e) \\ \mathcal{T}(v) \cdot \mathcal{T}(e') \\ \mathcal{T}(v) \cdot \mathcal{T}(e') \\ \mathcal{T}(e) \end{array} \right| \left| \begin{array}{l} s \\ \mathcal{T}(\langle \lambda.r, e' \rangle) \cdot s \\ \mathcal{T}(v) \cdot \mathcal{T}(\langle \lambda.r, e' \rangle) \cdot s \\ (c, \mathcal{T}(e)) \cdot s \\ \mathcal{T}(w) \cdot (c, \mathcal{T}(e)) \cdot s \\ \mathcal{T}(w) \cdot s \end{array} \right| \begin{array}{l} \text{by IH on } e \vdash t \Downarrow \langle \lambda.r, e' \rangle \\ \text{by IH on } e \vdash u \Downarrow v \\ \\ \text{by IH on } v \cdot e' \vdash r \Downarrow w \end{array}$$

Remark: this correctness property is called *partial*, because it applies only to terms t such that call-by-value evaluation $e \vdash t \Downarrow v$ succeeds. We could state a more precise property ensuring that other possible observable behaviours (non-termination, termination on error...) are also preserved, but the proof is significantly harder.

3 An abstract machine for call-by-name λ -calculus.

Call-by-name and call-by-need evaluation strategies do not use the same mechanisms than call-by-value. Machines for these strategies are thus also different.

Call-by-name, small-step semantics. The relation $t \rightarrow t'$ describes individual reduction steps. Call-by-name evaluation: apply a function to its argument without evaluating the argument first, then evaluate the argument on demand.

$$\frac{}{(\lambda x.t) u \rightarrow t\{x \leftarrow u\}} \qquad \frac{t \rightarrow t'}{t u \rightarrow t' u}$$

Differences with call-by-value reduction rules: the argument of a β -redex is not constrained anymore to be a value. Right members of applications are not evaluated. As for call-by-value, we will actually work with an environment rather than substitutions, and use the nameless representation based on de Bruijn indices.

The KAM (*Krivine Abstract Machine*) uses the same three kinds of elements than the SECD machine: a code pointer, an environment and a stack. Their contents is slightly different however.

- The environment still contains the actual parameters provided to a function, but these parameters are not evaluated anymore. Then the environment contains “terms to evaluate later” rather than values. These elements, called *thunks*, are pairs made of a code pointer and an environment (a thunk is similar to a functional closure, but its “code” part is not necessarily the code of a function).
- The stack contains parts of the term that are waiting to be considered, in particular the arguments to which a function has not been applied yet. Once again, these elements are thunks (and this contrasts with the stack of the SECD, which contained intermediate values).

The KAM has only three instructions. Compilation is as follows.

$$\begin{aligned} C(n) &= \text{ACCESS}(n) \\ C(\lambda.t) &= \text{GRAB}; C(t) \\ C(t u) &= \text{PUSH}(C(u)); C(t) \end{aligned}$$

- ACCESS(n): access the n -th thunk of the environment, and start evaluating it (in the SECD, the accessed element was rather pushed on the stack).
- PUSH(c): make a thunk with the code c and the current environment, and push it on the stack.
- GRAB: takes a thunk on the top of the stack, and move it to the environment.

Execution of a term $(\lambda.\lambda.t) a_1 a_2 a_3$ proceeds by:

1. pushing thunks for the arguments a_3 , a_2 and a_1 on the stack,
2. then applying the function $\lambda.\lambda.t$ by transferring the thunks for a_1 and a_2 from the stack to the environment,
3. finally executing the code of t , evaluating the thunks for a_1 or a_2 when the corresponding de Bruijn indices are accessed (and possibly transferring the thunk for a_3 to the environment if a new λ is encountered).

Machine steps. We write $\langle c, e \rangle$ for the thunk made with code c and environment e .

State before			State after			
code	env	stack	code	env	stack	
ACCESS(n); c	e	s	c'	e'	s	if $e[n] = \langle c', e' \rangle$
PUSH(c'); c	e	s	c	e	$\langle c', e \rangle \cdot s$	
GRAB; c	e	$\langle c', e' \rangle \cdot s$	c	$\langle c', e' \rangle \cdot e$	s	

The execution of this machine can be related step by step to small-step reduction of a λ -calculus extended with an explicit representation of the environment (*λ -calculus with explicit substitutions*).