



1 Definitional translation

We have seen how to encode a number of high-level language constructs—booleans, conditionals, natural numbers, and recursion—in λ -calculus. We now consider definitional translation, where we define the meaning of language constructs by translation to another language. This is a form of denotational semantics, but instead of the target being mathematical objects, it is a simpler programming language (such as λ -calculus). Note that definitional translation does not necessarily produce clean or efficient code; rather, it defines the meaning of the source language constructs in terms of the target language.

For each language construct, we will define an operational semantics directly, and then give an alternate semantics by translation to a simpler language. We will start by introducing *evaluation contexts*, which make it easier to present the new language features succinctly.

1.1 Evaluation contexts

Recall the syntax and CBV operational semantics for the lambda calculus:

$$\begin{aligned}
 e &::= x \mid \lambda x. e \mid e_1 e_2 \\
 v &::= \lambda x. e
 \end{aligned}$$

$$\frac{e_1 \rightarrow e'_1}{e_1 e_2 \rightarrow e'_1 e_2} \quad \frac{e_2 \rightarrow e'_2}{v_1 e_2 \rightarrow v_1 e'_2} \quad \beta\text{-REDUCTION} \frac{}{(\lambda x. e) v \rightarrow e\{v/x\}}$$

Of the operational semantics rules, only the β -reduction rule told us how to “reduce” an expression; the other two rules tell us the order to evaluate expressions—e.g., evaluate the left hand side of an application to a value first, then evaluate the right hand side of an application to a value. The operational semantics of many of the languages we will consider have this feature: there are two kinds of rules, *congruence rules* that specify evaluation order, and the *computation rules* that specify the “interesting” reductions.

Evaluation contexts are a simple mechanism that separates out these two kinds of rules. An evaluation context E (sometimes written $E[\cdot]$) is an expression with a “hole” in it, that is with a single occurrence of the special symbol $[\cdot]$ (called the “hole”) in place of a subexpression. Evaluation contexts are defined using a BNF grammar that is similar to the grammar used to define the language. The following grammar defines evaluation contexts for the pure CBV λ -calculus.

$$E ::= [\cdot] \mid E e \mid v E$$

We write $E[e]$ to mean the evaluation context E where the hole has been replaced with the expression e . The following are examples of evaluation contexts, and evaluation contexts with the

hole filled in by an expression.

$$\begin{aligned}
E_1 &= [\cdot] (\lambda x. x) & E_1[\lambda y. y y] &= (\lambda y. y y) \lambda x. x \\
E_2 &= (\lambda z. z z) [\cdot] & E_2[\lambda x. \lambda y. x] &= (\lambda z. z z) (\lambda x. \lambda y. x) \\
E_3 &= ([\cdot] \lambda x. x x) ((\lambda y. y) (\lambda y. y)) & E_3[\lambda f. \lambda g. f g] &= ((\lambda f. \lambda g. f g) \lambda x. x x) ((\lambda y. y) (\lambda y. y))
\end{aligned}$$

Using evaluation contexts, we can define the evaluation semantics for the pure CBV λ -calculus with just two rules, one for evaluation contexts, and one for β -reduction.

$$\frac{e \rightarrow e'}{E[e] \rightarrow E[e']} \qquad \beta\text{-REDUCTION} \frac{}{(\lambda x. e) v \rightarrow e\{v/x\}}$$

Note that the evaluation contexts for the CBV λ -calculus ensure that we evaluate the left hand side of an application to a value, and then evaluate the right hand side of an application to a value before applying β -reduction.

We can specify the operational semantics of CBN λ -calculus using evaluation contexts:

$$E ::= [\cdot] \mid E e \qquad \frac{e \rightarrow e'}{E[e] \rightarrow E[e']} \qquad \beta\text{-REDUCTION} \frac{}{(\lambda x. e_1) e_2 \rightarrow e_1\{e_2/x\}}$$

We'll see the benefit of evaluation contexts as we see languages with more syntactic constructs.

1.2 Multi-argument functions and currying

Our syntax for functions only allows function with a single argument: $\lambda x. e$. We could define a language that allows functions to have multiple arguments.

$$e ::= x \mid \lambda x_1, \dots, x_n. e \mid e_0 e_1 \dots e_n$$

Here, a function $\lambda x_1, \dots, x_n. e$ takes n arguments, with names x_1 through x_n . In a multi argument application $e_0 e_1 \dots e_n$, we expect e_0 to evaluate to an n -argument function, and e_1, \dots, e_n are the arguments that we will give the function.

We can define a CBV operational semantics for the multi-argument λ -calculus as follows.

$$E ::= [\cdot] \mid v_0 \dots v_{i-1} E e_{i+1} \dots e_n \qquad \frac{e \rightarrow e'}{E[e] \rightarrow E[e']}$$

$$\beta\text{-REDUCTION} \frac{}{(\lambda x_1, \dots, x_n. e_0) v_1 \dots v_n \rightarrow e_0\{v_1/x_1\}\{v_2/x_2\} \dots \{v_n/x_n\}}$$

The evaluation contexts ensure that we evaluate a multi-argument application $e_0 e_1 \dots e_n$ by evaluating each expression from left to right down to a value.

Now, the multi-argument λ -calculus isn't any more expressive than the pure λ -calculus. We can show this by showing how any multi-argument λ -calculus program can be translated into an equivalent pure λ -calculus program. We define a translation function $\mathcal{T}[\cdot]$ that takes an expression in the multi-argument λ -calculus and returns an equivalent expression in the pure

λ -calculus. That is, if e is a multi-argument lambda calculus expression, $\mathcal{T}[[e]]$ is a pure λ -calculus expression.

We define the translation as follows.

$$\begin{aligned}\mathcal{T}[[x]] &= x \\ \mathcal{T}[[\lambda x_1, \dots, x_n. e]] &= \lambda x_1. \dots \lambda x_n. \mathcal{T}[[e]] \\ \mathcal{T}[[e_0 e_1 e_2 \dots e_n]] &= (\dots ((\mathcal{T}[[e_0]] \mathcal{T}[[e_1]]) \mathcal{T}[[e_2]]) \dots \mathcal{T}[[e_n]])\end{aligned}$$

This process of rewriting a function that takes multiple arguments as a chain of functions that each take a single argument is called *currying*. Consider a mathematical function that takes two arguments, the first from domain A and the second from domain B , and returns a result from domain C . We could describe this function, using mathematical notation for domains of functions, as being an element of $A \times B \rightarrow C$. Currying this function produces a function that is an element of $A \rightarrow (B \rightarrow C)$. That is, the curried version of the function takes an argument from domain A , and returns a function that takes an argument from domain B and produces a result of domain C .

2 Products and let

A product is a pair of expressions (e_1, e_2) . If e_1 and e_2 are both values, then we regard the product as also being a value. (That is, we cannot further evaluate a product if both elements are values.) Given a product, we can access the first or second element using the operators $\#1$ and $\#2$ respectively. That is, $\#1 (v_1, v_2) \rightarrow v_1$ and $\#2 (v_1, v_2) \rightarrow v_2$. (Other common notation for projection includes π_1 and π_2 , and fst and snd .)

The syntax of λ -calculus extended with products and let expressions is defined as follows.

$$\begin{aligned}e ::= & x \mid \lambda x. e \mid e_1 e_2 \\ & \mid (e_1, e_2) \mid \#1 e \mid \#2 e \\ & \mid \text{let } x = e_1 \text{ in } e_2 \\ v ::= & \lambda x. e \mid (v_1, v_2)\end{aligned}$$

Note that values in this language are either functions or pairs of values.

We define a small-step CBV operational semantics for the language using evaluation contexts.

$$\begin{aligned}E ::= & [\cdot] \mid E e \mid v E \mid (E, e) \mid (v, E) \mid \#1 E \mid \#2 E \mid \text{let } x = E \text{ in } e_2 \\ \frac{e \rightarrow e'}{E[e] \rightarrow E[e']} & \qquad \beta\text{-REDUCTION} \frac{}{(\lambda x. e) v \rightarrow e\{v/x\}} \\ \frac{}{\#1 (v_1, v_2) \rightarrow v_1} & \qquad \frac{}{\#2 (v_1, v_2) \rightarrow v_2} \\ \frac{}{\text{let } x = v \text{ in } e \rightarrow e\{v/x\}} & \end{aligned}$$

Next, we define an equivalent semantics by translation to the pure CBV λ -calculus.

$$\begin{aligned}
\mathcal{T}[[x]] &= x \\
\mathcal{T}[[\lambda x. e]] &= \lambda x. \mathcal{T}[[e]] \\
\mathcal{T}[[e_1 e_2]] &= \mathcal{T}[[e_1]] \mathcal{T}[[e_2]] \\
\mathcal{T}[[e_1, e_2]] &= (\lambda x. \lambda y. \lambda f. f x y) \mathcal{T}[[e_1]] \mathcal{T}[[e_2]] \\
\mathcal{T}[[\#1 e]] &= \mathcal{T}[[e]] (\lambda x. \lambda y. x) \\
\mathcal{T}[[\#2 e]] &= \mathcal{T}[[e]] (\lambda x. \lambda y. y) \\
\mathcal{T}[[\text{let } x = e_1 \text{ in } e_2]] &= (\lambda x. \mathcal{T}[[e_2]]) \mathcal{T}[[e_1]]
\end{aligned}$$

Note that we encode a pair (e_1, e_2) as a value that takes a function f , and applies f to v_1 and v_2 , where v_1 and v_2 are the result of evaluating e_1 and e_2 respectively. The projection operators pass a function to the encoding of pairs that selects either the first or second element as appropriate. Also note that the expression $\text{let } x = e_1 \text{ in } e_2$ is equivalent to the application $(\lambda x. e_2) e_1$.

3 Laziness

In previous lectures we defined semantics for both the call-by-name λ -calculus and the call-by-value λ -calculus. It turns out that we can translate a call-by-name program into a call-by-value program. In CBV, arguments to functions are evaluated before the function is applied; in CBN, functions are applied as soon as possible. In the translation, we delay the evaluation of arguments by wrapping them in a function. This is called a *thunk*: wrapping a computation in a function to delay its evaluation.

Since arguments to functions are turned into thunks, when we want to use an argument in a function body, we need to evaluate the thunk. We do so by applying the thunk (which is simply a function); it doesn't matter what we apply the thunk to, since the thunk's argument is never used.

$$\begin{aligned}
\mathcal{T}[[x]] &= x (\lambda y. y) \\
\mathcal{T}[[\lambda x. e]] &= \lambda x. \mathcal{T}[[e]] \\
\mathcal{T}[[e_1 e_2]] &= \mathcal{T}[[e_1]] (\lambda z. \mathcal{T}[[e_2]]) \quad z \text{ is not a free variable of } e_2
\end{aligned}$$

4 References

We can also introduce constructs for creating, reading, and updating memory locations, also called *references*. The resulting language is still a functional language (since functions are first-class values), but expressions can have side-effects, that is, they can modify state. The syntax of this language is defined as follows.

$$\begin{aligned}
e &::= x \mid \lambda x. e \mid e_0 e_1 \mid \text{ref } e \mid !e \mid e_1 := e_2 \mid \ell \\
v &::= \lambda x. e \mid \ell
\end{aligned}$$

Expression `ref e` creates a new memory location (like a `malloc`), and sets the initial contents of the location to (the result of) `e`. The expression `ref e` itself evaluates to a memory location ℓ . Think of a location as being like a pointer to a memory address. The expression `!e` assumes that `e` evaluates to a memory location, and `!e` evaluates to the current contents of the memory location. Expression `e1 := e2` assumes that `e1` evaluates to a memory location ℓ , and updates the contents of ℓ with (the result of) `e2`. Locations ℓ are not intended to be used directly by a programmer: they are not part of the *surface syntax* of the language, the syntax that a programmer would write. They are introduced only by the operational semantics.

We define a small-step CBV operational semantics. We use configurations $\langle \sigma, e \rangle$, where `e` is an expression, and σ is a map from locations to values.

$$\begin{array}{c}
 E ::= [\cdot] \mid E e \mid v E \mid \text{ref } E \mid !E \mid E := e \mid v := E \\
 \beta\text{-REDUCTION} \frac{}{\langle \sigma, (\lambda x. e) v \rangle \rightarrow \langle \sigma, e\{v/x\} \rangle} \quad \text{ALLOC} \frac{}{\langle \sigma, \text{ref } v \rangle \rightarrow \langle \sigma[\ell \mapsto v], \ell \rangle} \ell \notin \text{dom}(\sigma) \\
 \text{DEREF} \frac{}{\langle \sigma, !\ell \rangle \rightarrow \langle \sigma, v \rangle} \sigma(\ell) = v \quad \text{ASSIGN} \frac{}{\langle \sigma, \ell := v \rangle \rightarrow \langle \sigma[\ell \mapsto v], v \rangle}
 \end{array}$$

References do not add any expressive power to the λ -calculus: it is possible to translate λ -calculus with references to the pure λ -calculus. Intuitively, this is achieved by explicitly representing the store, and threading the store through the evaluation of the program. The details are left as an exercise.

5 Adequacy of translation

In each of the previous translations, we defined a semantics for the source language (using evaluation contexts and small-step rules) and the target language (by translation). We would like to be able to show that the translation is correct—that is, that it preserves the meaning of source programs.

More precisely, we would like an expression `e` in the source language to evaluate to a value `v` if and only if the translation of `e` evaluates to a value `v'` such that `v'` is “equivalent to” `v`. What exactly it means for `v'` to be “equivalent to” `v` will depend on the translation. Sometimes, it will mean that `v'` is literally the translation of `v`; other times, it will mean that `v'` is merely related to the translation of `v` by some equivalence.

One tricky issue is that in general, there can be many ways to define equivalences on functions. One way is to say that two functions are equivalent if they agree on the result when applied to any value of a base type (e.g., integers or booleans). The idea is that if two functions disagree when passed a more complex value (say, a function), then we could write a program that uses these functions to produce functions that disagree on values of base types.

There are two criteria for a translation to be *adequate*: soundness and completeness. For clarity, let’s suppose that Exp_{src} is the set of source language expressions, and that \rightarrow_{src} and \rightarrow_{trg} are the evaluation relations for the source and target languages respectively. A translation is sound if

every target evaluation represents a source evaluation:

Soundness: $\forall e \in \mathbf{Exp}_{\text{src}}. \text{if } \mathcal{T}[[e]] \rightarrow_{\text{trg}}^* v' \text{ then } \exists v. e \rightarrow_{\text{src}}^* v \text{ and } v' \text{ equivalent to } v$

A translation is complete if every source evaluation has a target evaluation.

Completeness: $\forall e \in \mathbf{Exp}_{\text{src}}. \text{if } e \rightarrow_{\text{src}}^* v \text{ then } \exists v'. \mathcal{T}[[e]] \rightarrow_{\text{trg}}^* v' \text{ and } v' \text{ equivalent to } v$