

Lecture Notes on The Lambda Calculus

15-814: Types and Programming Languages
Frank Pfenning

Lecture 1
Tuesday, September 1, 2020

1 Introduction

This course is about the principles of programming language design, many of which derive from the notion of *type*. Nevertheless, we will start by studying an exceedingly pure notion of computation based only on the notion of function, that is, Church's λ -calculus [CR36]. There are several reasons to do so.

- We will see a number of important concepts in their simplest possible form, which means we can discuss them in full detail. We will then reuse these notions frequently throughout the course without the same level of detail.
- The λ -calculus is of great historical and foundational significance. The independent and nearly simultaneous development of Turing Machines [Tur36] and the λ -Calculus [CR36] as universal computational mechanisms led to the *Church-Turing Thesis*, which states that the effectively computable (partial) functions are exactly those that can be implemented by Turing Machines or, equivalently, in the λ -calculus.
- The notion of function is the most basic abstraction present in nearly all programming languages. If we are to study programming languages, we therefore must strive to understand the notion of function.
- It's cool!

2 The λ -Calculus

In ordinary mathematical practice, functions are ubiquitous. For example, we might define

$$\begin{aligned} f(x) &= x + 5 \\ g(y) &= 2 * y + 7 \end{aligned}$$

Oddly, we never state what f or g actually are, we only state what happens when we apply them to arbitrary arguments such as x or y . The λ -calculus starts with the simple idea that we should have notation for the function itself, the so-called λ -abstraction.

$$\begin{aligned} f &= \lambda x. x + 5 \\ g &= \lambda y. 2 * y + 7 \end{aligned}$$

In general, $\lambda x. e$ for some arbitrary expression e stands for the function which, when applied to some e' becomes $[e'/x]e$, that is, the result of *substituting* or *plugging in* e' for occurrences of the variable x in e . For now, we will use this notion of substitution informally—in the next lecture we will define it formally.

We can already see that in a pure calculus of functions we will need at least three different kinds of expressions: λ -abstractions $\lambda x. e$ to form function, *application* $e_1 e_2$ to apply a function e_1 to an argument e_2 , and *variables* x, y, z , etc. We summarize this in the following form

$$\begin{array}{ll} \text{Variables} & x \\ \text{Expressions } e & ::= \lambda x. e \mid e_1 e_2 \mid x \end{array}$$

This is not the definition of the *concrete syntax* of a programming language, but a slightly more abstract form called *abstract syntax*. When we write down concrete expressions there are additional conventions and notations such as parentheses to avoid ambiguity.

1. Juxtaposition (which expresses application) is *left-associative* so that $x y z$ is read as $(x y) z$
2. $\lambda x.$ is a prefix whose scope extends as far as possible while remaining consistent with the parentheses that are present. For example, $\lambda x. (\lambda y. x y z) x$ is read as $\lambda x. ((\lambda y. (x y) z) x)$.

We say $\lambda x. e$ *binds* the variable x with scope e . Variables that occur in e but are not bound are called *free variables*, and we say that a variable x may occur free in an expression e . For example, y is free in $\lambda x. x y$ but not

x . Bound variables can be renamed consistently in a term So $\lambda x. x + 5 = \lambda y. y + 5 = \lambda \textit{whatever}. \textit{whatever} + 5$. Generally, we rename variables *silently* because we identify terms that differ only in the names of λ -bound variables. But, if we want to make the step explicit, we call it α -conversion.

$$\lambda x. e =_{\alpha} \lambda y. [y/x]e \quad \text{provided } y \text{ not free in } e$$

The proviso is necessary, for example, because $\lambda x. x y \neq \lambda y. y y$.

We capture the rule for function application with

$$(\lambda x. e_2) e_1 =_{\beta} [e_1/x]e_2$$

and call it β -conversion. Some care has to be taken for the substitution to be carried out correctly—we will return to this point later.

If we think beyond mere equality at *computation*, we see that β -conversion has a definitive direction: we apply it from left to right. We call this β -reduction and it is the engine of computation in the λ -calculus.

$$(\lambda x. e_2) e_1 \longrightarrow_{\beta} [e_1/x]e_2$$

3 Simple Functions and Combinators

The simplest functions are the identity function and the constant function. The identity function, called I , just returns its argument x .

$$I = \lambda x. x$$

The constant function returning x could be written as

$$\lambda y. x$$

We calculate

$$(\lambda y. x) e \longrightarrow_{\beta} x$$

for any expression e since y does not occur in the expression x . This is somewhat incomplete in the sense the expression $\lambda y. x$ has a *free variable* which is therefore fixed. What we would like is a *closed expression* K (one without free variables) such $K x$ is the constant function, always returning x . But that's easy: we just abstract over x !

$$K = \lambda x. \lambda y. x$$

Then $K x \longrightarrow_{\beta} \lambda y. x$ is the constant function returning x .

A combinator for us is just a closed λ -expression like I or K . We will see more interesting combinators in the next lecture.

4 Summary of λ -Calculus

λ -Expressions.

Variables x
 Expressions $e ::= \lambda x. e \mid e_1 e_2 \mid x$

$\lambda x. e$ binds x with scope e , which is as large as possible while remaining consistent with the given parentheses. Juxtaposition $e_1 e_2$ is left-associative.

Equality.

Substitution $[e_1/x]e_2$ (capture-avoiding, see Lecture 2)
 α -conversion $\lambda x. e =_\alpha \lambda y. [y/x]e$ provided y not free in e
 β -conversion $(\lambda x. e_2) e_1 =_\beta [e_1/x]e_2$

We generally apply α -conversion silently, identifying terms that differ only in the names of the bound variables.

Reduction.

β -reduction $(\lambda x. e_2) e_1 \longrightarrow_\beta [e_1/x]e_2$

5 Representing Booleans

Before we can claim the λ -calculus as a universal language for computation, we need to be able to represent *data*. The simplest nontrivial data type are the Booleans, a type with two elements: *true* and *false*. The general technique is to represent the values of a given type by *normal forms*, that is, expressions that cannot be reduced. Furthermore, they should be *closed*, that is, not contain any free variables. We need to be able to distinguish between two values, and in a closed expression that suggest introducing two bound variables. We then define rather arbitrarily one to be *true* and the other to be *false*

$true = \lambda x. \lambda y. x$
 $false = \lambda x. \lambda y. y$

The next step will be to define *functions* on values of the type. Let's start with negation: we are trying to define a λ -expression *not* such that

$not\ true =_\beta\ false$
 $not\ false =_\beta\ true$

We start with the obvious:

$$\text{not} = \lambda b. \dots$$

Now there are two possibilities: we could either try to apply b to some arguments, or we could build some λ -abstractions. In lecture, we followed both paths. Let's first try the one where b is applied to some arguments.

$$\text{not} = \lambda b. b (\dots) (\dots)$$

We suggest two arguments to b , because b stands for a Boolean, and Booleans true and false both take two arguments. $\text{true} = \lambda x. \lambda y. x$ will pick out the first of these two arguments and discard the second, so since we specified $\text{not true} = \text{false}$, the first argument to b should be false !

$$\text{not} = \lambda b. b \text{false} (\dots)$$

Since $\text{false} = \lambda x. \lambda y. y$ picks out the second argument and $\text{not false} = \text{true}$, the second argument to b should be true .

$$\text{not} = \lambda b. b \text{false true}$$

Now it is a simple matter to calculate that the computation of not applied to true or false completes in three steps and obtain the correct result.

$$\begin{array}{l} \text{not true} \quad \longrightarrow_{\beta}^3 \text{false} \\ \text{not false} \quad \longrightarrow_{\beta}^3 \text{true} \end{array}$$

We write $\longrightarrow_{\beta}^n$ for reduction in n steps, and $\longrightarrow_{\beta}^*$ for reduction in an arbitrary number of steps, including zero steps. In other words, $\longrightarrow_{\beta}^*$ is the reflexive and transitive closure of \longrightarrow_{β} .

An alternative solution hinted at above is to start with

$$\text{not}' = \lambda b. \lambda x. \lambda y. \dots$$

We pose this because the result of $\text{not } b$ should be a Boolean, and the two Booleans both start with two λ -abstractions. Now we reuse the previous idea, but apply b not to false and true , but to y and x .

$$\text{not}' = \lambda b. \lambda x. \lambda y. b y x$$

Again, we calculate

$$\begin{array}{l} \text{not}' \text{true} \quad \longrightarrow_{\beta}^3 \text{false} \\ \text{not}' \text{false} \quad \longrightarrow_{\beta}^3 \text{true} \end{array}$$

An important observation here is that

$$\text{not} = \lambda b. b (\lambda x. \lambda y. y) (\lambda x. \lambda y. x) \neq \lambda b. \lambda x. \lambda y. b y x = \text{not}'$$

Both of these are *normal forms* (they cannot be reduced) and therefore represent *values* (the results of computation). Both correctly implement negation on Booleans, but they are *different*. This is evidence that when computing with particular data representations in the λ -calculus it is *not extensional*: even though the functions behave the same on all the arguments we care about (here just *true* and *false*), they are not convertible. To actually see that they are not convertible we need the Church-Rosser theorem which says if e_1 and e_2 are $\alpha\beta$ -convertible then there is a common reduct e such that $e_1 \rightarrow_{\beta}^* e$ and $e_2 \rightarrow_{\beta}^* e$.

As a next exercise we try exclusive conjunction. We want to define a λ -expression *and* such that

$$\begin{aligned} \text{and } \text{true } \text{true} &=_{\beta} \text{true} \\ \text{and } \text{true } \text{false} &=_{\beta} \text{false} \\ \text{and } \text{false } \text{true} &=_{\beta} \text{false} \\ \text{and } \text{false } \text{false} &=_{\beta} \text{false} \end{aligned}$$

Learning from the negation, we start by guessing

$$\text{and} = \lambda b. \lambda c. b (\dots) (\dots)$$

where we arbitrarily put b first. Looking at the equations, we see that if b is *true* then the result is always c .

$$\text{and} = \lambda b. \lambda c. b c (\dots)$$

If b is *false* the result is always just *false*, no matter what c is.

$$\text{and} = \lambda b. \lambda c. b c \text{false}$$

Again, it is now a simple matter to verify the desired equations and that, in fact, the right-hand side of these equations is obtained by reduction.

6 The LAMBDA Language

In lecture, we used a toy implementation of the λ -calculus in a language called LAMBDA. This implementation uses a *concrete syntax* where λ is

written as a backslash '\'. A program consists of a sequence of declarations, of which there are three forms:

defn $x = e$ variable x stands for e
norm $x = e$ variable x stands for the normal form of e
conv $e_1 = e_2$ verify that e_1 and e_2 have the same normal form

Allowing definitions is a convenience, but it does not change the expressive power of the λ -calculus, because we can replace **defn** $x = e$ by $(\lambda x. \dots) e$ where ' \dots ' represents the scope of the definition. The **norm** and **conv** declarations initiate computation and allow the programmer to examine the normal form of an expression (if it exists).

In addition, declarations can be negated with **!**, for example, to check that two expressions are not convertible.

```

1  % represent booleans as closed expressions in normal form
2
3  defn true = \x. \y. x
4  defn false = \x. \y. y
5
6  defn not = \b. b false true
7  defn not' = \b. \x. \y. b y x
8
9  (* confirm that not and not' are not convertible *)
10 !conv not = not'
11
12 % normalize "not true"
13 norm _ = not true
14
15 % test not and not' against their specification
16 conv not true = false
17 conv not false = true
18
19 conv not' true = false
20 conv not' false = true

```

Listing 1: Booleans in LAMBDA

For more information on LAMBDA, consult the [Software](#) page for the course.

Exercises

Exercise 1 Define the following functions on Booleans in at least two distinct ways.

1. Exclusive or “*xor*”.
2. The conditional “*if*” such that

$$\begin{aligned} \text{if true } e_1 e_2 &=_{\beta} e_1 \\ \text{if false } e_1 e_2 &=_{\beta} e_2 \end{aligned}$$

References

- [CR36] Alonzo Church and J.B. Rosser. Some properties of conversion. *Transactions of the American Mathematical Society*, 39(3):472–482, May 1936.
- [Tur36] Alan Turing. On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42:230–265, 1936. Published 1937.