A Tutorial Introduction to the Lambda Calculus

Raúl Rojas[*]

FU Berlin, WS-97/98

Abstract

This paper is a short and painless introduction to the λ calculus. Originally developed in order to study some mathematical properties of effectively computable functions, this formalism has provided a strong theoretical foundation for the family of functional programming languages. We show how to perform some arithmetical computations using the λ calculus and how to define recursive functions, even though functions in λ calculus are not given names and thus cannot refer explicitly to themselves.

1 Definition

The λ calculus can be called the *smallest universal programming language of the world.* The λ calculus consists of a single transformation rule (variable substitution) and a single function definition scheme. It was introduced in the 1930s by Alonzo Church as a way of formalizing the concept of effective computability. The λ calculus is universal in the sense that any computable function can be expressed and evaluated using this formalism. It is thus equivalent to Turing machines. However, the λ calculus emphasizes the use of transformation rules and does not care about the actual machine implementing them. It is an approach more related to software than to hardware.

The central concept in λ calculus is the "expression". A "name", also called a "variable", is an identifier which, for our purposes, can be any of the letters $a, b, c, \ldots$ An expression is defined recursively as follows:

$$
\begin{aligned}
\text{<expression>} \quad &:= \quad \text{<name> | <function> | <application>} \\
\text{<function>} \quad &:= \quad \lambda \text{ <name>.<expression>} \\
\text{<application>} \quad &:= \quad \text{<expression><expression>}
\end{aligned}
$$

An expression can be surrounded with parenthesis for clarity, that is, if E is an expression, (E) is the same expression. The only keywords used in the language are λ and the dot. In order to avoid cluttering expressions with parenthesis, we adopt the convention that function application associates from the left, that is, the expression

$$E_1 E_2 E_3 \ldots E_n$$

is evaluated applying the expressions as follows:

$$(\ldots((E_1 E_2)E_3)\ldots E_n)$$

As can be seen from the definition of λ expressions given above, a single identifier is a λ expression. An example of a function is the following:

$$\lambda x.x$$

This expression defines the identity function. The name after the λ is the identifier of the argument of this function. The expression after the point (in this case a single x) is called the "body" of the definition.

Functions can be applied to expressions. An example of an application is

$$(\lambda x.x)y$$

This is the identity function applied to y. Parenthesis are used for clarity in order to avoid ambiguity. Function applications are evaluated by substituting the value of the argument x (in this case y) in the body of the function definition, i.e.

$$(\lambda x.x)y = [y/x]x = y$$

In this transformation the notation $[y/x]$ is used to indicate that all occurrences of x are substituted by y in the expression to the right.

The names of the arguments in function definitions do not carry any meaning by themselves. They are just "place holders", that is, they are used to indicate how to rearrange the arguments of the function when it is evaluated. Therefore

$$(\lambda z.z) \equiv (\lambda y.y) \equiv (\lambda t.t) \equiv (\lambda u.u)$$

and so forth. We use the symbol "$\equiv$" to indicate that when $A \equiv B$, A is just a synonym of B.

1.1 Free and bound variables

In λ calculus all names are local to definitions. In the function $\lambda x.x$ we say that x is "bound" since its occurrence in the body of the definition is preceded by λx. A name not preceded by a λ is called a "free variable". In the expression

$$(\lambda x.xy)$$

the variable x is bound and y is free. In the expression

$$(\lambda x.x)(\lambda y.yx)$$

the x in the body of the first expression from the left is bound to the first λ. The y in the body of the second expression is bound to the second λ and the x is free. It is very important to notice that the x in the second expression is totally independent of the x in the first expression.

Formally we say that a variable <name> is free in an expression if one of the following three cases holds:

- <name> is free in <name>.

- <name> is free in λ <name$_1$ > . <exp> if the identifier <name>$\neq$<name$_1$ > and <name> is free in <exp>.

- <name> is free in $E_1 E_2$ if <name> is free in E_1 or if it is free in E_2.

A variable <name> is bound if one of two cases holds:

- <name> is bound in λ <name$_1$ > . <exp> if the identifier <name>=<name$_1$ > or if <name> is bound in <exp>.

- <name> is bound in $E_1 E_2$ if <name> is bound in E_1 or if it is bound in E_2.

It should be emphasized that the same identifier can occur free and bound in the same expression. In the expression

$$(\lambda x.xy)(\lambda y.y)$$

the first y is free in the parenthesized subexpression to the left. It is bound in the subexpression to the right. It occurs therefore free as well as bound in the whole expression.

1.2 Substitutions

The more confusing part of standard λ calculus, when first approaching it, is the fact that we do not give names to functions. Any time we want to apply a function, we write the whole function definition and then procede to evaluate it. To simplify the notation, however, we will use capital letters, digits and other symbols as synonyms for some function definitions. The identity function, for example, can be denoted with I which is a synonym for $(\lambda x.x)$.

The identity function applied to itself is the application

$$\mathsf{II} \equiv (\lambda x.x)(\lambda x.x)$$

In this expression the first x in the body of the first expression in parenthesis is independent of the x in the body of the second expression. We can in fact rewrite the above expression as

$$\mathsf{II} \equiv (\lambda x.x)(\lambda z.z)$$

The identity function applied to itself

$$\mathsf{II} \equiv (\lambda x.x)(\lambda z.z)$$

yields therefore

$$[\lambda z.z/x]x = \lambda z.z \equiv \mathsf{I}$$

that is, the identity function again.

We should be careful when performing substitutions to avoid mixing up free occurrences of an identifier with bound ones. In the expression

$$(\lambda x.(\lambda y.xy))y$$

the function to the left contains a bound y, whereas the y at the right is free. An incorrect substitution would mix the two identifiers in the erroneous result

$$(\lambda y.yy)$$

Simply by renaming the bound y to t we obtain

$$(\lambda x.(\lambda t.xt))y = (\lambda t.yt)$$

which is a completely different result but nevertheless the correct one.

Therefore, if the function $\lambda x. <exp>$ is applied to E, we substitute all *free* occurrences of x in $<exp>$ with E. If the substitution would bring a free variable of E in an expression where this variable occurs bound, we rename the bound variable before performing the substitution. For example, in the expression

$$(\lambda x.(\lambda y.(x(\lambda x.xy))))y$$

we associate the argument x with y. In the body

$$(\lambda y.(x(\lambda x.xy)))$$

only the first x is free and can be substituted. Before substituting though, we have to rename the variable y to avoid mixing its bound with is free occurrence:

$$[y/x](\lambda t.(x(\lambda x.xt))) = (\lambda t.(y(\lambda x.xt)))$$

In normal order reduction we try to reduce always the left most expression of a series of applications. We continue until no further reductions are possible.

2 Arithmetic

We expect from a programming language that it should be capable of doing arithmetical calculations. Numbers can be represented in lambda calculus starting from zero and writing "suc(zero)" to represent 1, "suc(suc(zero))" to represent 2, and so on. In the lambda calculus we can only define new functions. Numbers will be defined as functions using the following approach: zero can be defined as

$$\lambda s.(\lambda z.z)$$

This is a function of two arguments s and z. We will abbreviate such expressions with more than one argument as

$$\lambda sz.z$$

It is understood here that s is the first argument to be substituted during the evaluation and z the second. Using this notation, the first natural numbers can be defined as

$$
\begin{array}{rcl}
1 & \equiv & \lambda sz.s(z) \\
2 & \equiv & \lambda sz.s(s(z)) \\
3 & \equiv & \lambda sz.s(s(s(z)))
\end{array}
$$

and so on.

Our first interesting function is the successor function. This can be defined as

$$\mathsf{S} \equiv \lambda wyx.y(wyx)$$

The successor function applied to our representation for zero yields

$$\mathsf{S0} \equiv (\lambda wyx.y(wyx))(\lambda sz.z)$$

In the body of the first expression we substitute all occurrences of w with $(\lambda sz.z)$ and this yields

$$\lambda yx.y((\lambda sz.z)yx) = \lambda yx.y((\lambda z.z)x) = \lambda yx.y(x) \equiv \mathsf{1}$$

That is, we obtain the representation of the number $\mathsf{1}$ (remember that variable names are "dummies").

Successor applied to 1 yields:

$$\mathsf{S1} \equiv (\lambda wyx.y(wyx))(\lambda sz.s(z)) = \lambda yx.y((\lambda sz.s(z))yx) = \lambda yx.y(y(x)) \equiv \mathsf{2}$$

Notice that the only purpose of applying the number $(\lambda sz.s(z))$ to the arguments y and x is to "rename" the variables used in the definition of our number.

2.1 Addition

Addition can be obtained immediately by noting that the body sz of our definition of the number 1, for example, can be interpreted as the application of the function s on z. If we want to add say 2 and 3, we just apply the successor function two times to 3.

Let us try the following in order to compute 2+3:

$$\mathsf{2S3} \equiv (\lambda sz.s(sz))(\lambda wyx.y(wyx))(\lambda uv.u(u(uv)))$$

The first expression on the right side is a 2, the second is the successor function, the third is a 3 (we have renamed the variables for clarity). The expression above reduces to

$$(\lambda wyx.y((wy)x))((\lambda wyx.y((wy)x))(\lambda uv.u(u(uv)))) \equiv \mathsf{SS3}$$

The reader can verify that $\mathsf{SS3}$ reduces to $\mathsf{S4} = \mathsf{5}$.

2.2 Multiplication

The multiplication of two numbers x and y can be computed using the following function:

$$(\lambda xyz.x(yz))$$

The product of 2 by 2 is then:

$$(\lambda xyz.x(yz))\mathsf{22}$$

which reduces to

$$(\lambda z.\mathsf{2}(\mathsf{2}z))$$

The reader can verify that by further reducing this expression, we can obtain the expected result 4.

3 Conditionals

We introduce the following two functions which we call the values "true"

$$\mathsf{T} \equiv \lambda xy.x$$

and "false"

$$\mathsf{F} \equiv \lambda xy.y$$

The first function takes two arguments and returns the first one, the second function returns the second of two arguments.

3.1 Logical operations

It is now possible to define logical operations using this representation of the truth values.

The AND function of two arguments can be defined as

$$\wedge \equiv \lambda xy.xy(\lambda uv.v) \equiv \lambda xy.xy\mathsf{F}$$

The OR function of two arguments can be defined as

$$\vee \equiv \lambda xy.x(\lambda uv.u)y \equiv \lambda xy.x\mathsf{T}y$$

Negation of one argument can be defined as

$$\neg \equiv \lambda x.x(\lambda uv.v)(\lambda ab.a) \equiv \lambda x.x\mathsf{FT}$$

The negation function applied to "true" is

$$\neg\mathsf{T} \equiv \lambda x.x(\lambda uv.v)(\lambda ab.a)(\lambda cd.c)$$

which reduces to

$$\mathsf{TFT} \equiv (\lambda cd.c)(\lambda uv.v)(\lambda ab.a) = (\lambda uv.v) \equiv \mathsf{F}$$

that is, the truth value "false".

3.2 A conditional test

It is very convenient in a programming language to have a function which is true if a number is zero and false otherwise. The following function Z complies with this requirement

$$\mathsf{Z} \equiv \lambda x.x\mathsf{F}\neg\mathsf{F}$$

To understand how this function works, note that

$$0fa \equiv (\lambda sz.z)fa = a$$

that is, the function f applied zero times to the argument a yields a. On the other hand, F applied to any argument yields the identity function

$$\mathsf{F}a \equiv (\lambda xy.y)a = \lambda y.y \equiv \mathsf{I}$$

6

We can now test if the function Z works correctly. The function applied to zero yields

$$\mathsf{Z}0 \equiv (\lambda x.x\mathsf{F}\neg\mathsf{F})0 = 0\mathsf{F}\neg\mathsf{F} = \neg\mathsf{F} = \mathsf{T}$$

because F applied 0 times to $\neg$ yields $\neg$. The function Z applied to any other number N yields

$$\mathsf{Z}\mathsf{N} \equiv (\lambda x.x\mathsf{F}\neg\mathsf{F})\mathsf{N} = \mathsf{N}\mathsf{F}\neg\mathsf{F}$$

The function F is then applied N times to $\neg$. But F applied to anything is the identity, so that the above expression reduces for any number N greater than zero to

$$\mathsf{I}\mathsf{F} = \mathsf{F}$$

3.3 The predecessor function

We can now define the predecessor function combining some of the functions introduced above. When looking for the predecessor of n, the general strategy will be to create a pair $(n, n-1)$ and then pick the second element of the pair as the result.

A pair (a, b) can be represented in λ-calculus using the function

$$(\lambda z.zab)$$

We can extract the first element of the pair from the expression applying this function to T

$$(\lambda z.zab)\mathsf{T} = \mathsf{T}ab = a$$

and the second applying the function to F

$$(\lambda z.zab)\mathsf{F} = \mathsf{F}ab = b$$

The following function generates from the pair $(n, n-1)$ (which is the argument p in the function) the pair $(n+1, n-1)$:

$$\Phi \equiv (\lambda pz.z(\mathsf{S}(p\mathsf{T}))(p\mathsf{T}))$$

The subexpression $p\mathsf{T}$ extracts the first element from the pair p. A new pair is formed using this element, which is incremented for the first position of the new pair and just copied for the second position of the new pair.

The predecessor of a number n is obtained by applying n times the function Φ to the pair $(\lambda.z00)$ and then selecting the second member of the new pair:

$$\mathsf{P} \equiv (\lambda n.n\Phi(\lambda z.z00)\mathsf{F}$$

Notice that using this approach the predecessor of zero is zero. This property is useful for the definition of other functions.

3.4 Equality and inequalities

With the predecessor function as the building block, we can now define a function which tests if a number x is greater than or equal to a number y:

$$\mathsf{G} \equiv (\lambda xy.\mathsf{Z}(x\mathsf{P}y))$$