

Programming Language Theory

Lecture notes

5 Lambda calculus

Lambda calculus is a formal system based on a notation for functions that was introduced by Church in 1930. The objective was to make a general theory of functions and to provide a formal basis for logic and mathematics.

The first goal has established a rich and beautiful theory which has inspired the construction of functional programming languages such as Lisp and Haskell and has been used to define formal semantics for programming languages.

The objects of the theory are *terms* which may be interpreted both as functions defined by computational rules and values which are the arguments of the functions. There is no distinction between functions and values and a function may even be applied to itself.

There are terms called *combinators* in the λ -calculus with a certain property which can describe computations without variables. The corresponding theory is called *combinatory logic* having Schönfinkel and Curry as originators.

Lambda notation

A standard way to define a function in mathematics is to give it a name and give the value of the function as an expression in the formal arguments the function. A simple example is the successor function, $f(x) = x + 1$. Using λ -notation we can express this function as a λ -expression without giving it a name. We write $\lambda x . x + 1$. Just as f may be applied to the argument 1, $f(1)$, we may apply the lambda expression to the same argument, $(\lambda x . x + 1)(1)$. It is customary to omit the parentheses around the argument when it is a simple constant or variable, so we write $(\lambda x . x + 1) 1$.

We evaluate both $f(1)$ and $(\lambda x . x + 1) 1$ by substituting 1 for x in $x + 1$ and performing the addition.

In mathematics a function of two variables may be defined like $f(x, y) = x + y$. With λ -notation we would write $\lambda(x, y) . x + y$ to denote the same function. There is, however, another way to regard the addition function, not usually done in mathematics. Using lambda notation we could write $\lambda x . \lambda y . x + y$ or more clearly $\lambda x . (\lambda y . x + y)$. Applying this λ -expression to 1 yields $\lambda y . 1 + y$ after substituting 1 for x . So the result of the application is a function, the successor function. Starting with $((\lambda x . \lambda y . x + y) 1) 2$ we get $(\lambda y . 2 + y) 1$ and $2 + 1$ after two substitutions.

In mathematics there are also functions returning functions as values and having function arguments. They are usually called *operators* or *functionals*. A well known example is the differentiation operator $\frac{d}{dx} x^2 = 2x$. Another one is composition of functions. If f and g are functions, e.g. from R to R , $h = f \circ g$ is defined by $h(x) = f(g(x))$. The differentiation and composition operators are functions according to the standard set theoretic definition of a function as a set of pairs.

Grammar for the lambda calculus

The objects of the λ -calculus are called λ -terms. There are three kinds of λ -terms. The simplest one is an *identifier*. The form of an identifier is not very important. The theory needs an unbounded supply of identifiers, but in our examples we will usually manage with a small number of them. We could allow Java identifiers, but it is customary to use single letter identifiers, possibly with an index. Examples: x, y, z, x_0 and x_1 .

The second kind is an *abstraction*. An abstraction can be thought of as a function with a formal parameter and a function body which is a λ -term. A simple example is the *identity* abstraction $\lambda x. x$. When an abstraction is part of a more complex λ -term we will enclose it in parentheses, $(\lambda x. x)$.

The third kind is an *application*. It can be thought of as an application of a function to its argument, but any λ -term may be applied to any λ -term, so $x y$ is a legal application as is $(\lambda x. x) y$. When an application appears inside a complex term it may be surrounded by parentheses, $((\lambda x. x) y)$.

The first concrete grammar requires a parentheses pair around each abstraction and each application. Using Backus-Naur formalism we write

$$\textit{term} \rightarrow \textit{id} \mid (\lambda \textit{id} . \textit{term}) \mid (\textit{term} \textit{term})$$

where a *term* is defined to be an identifier, *id*, an abstraction with an identifier and a *term* or an application with two *terms*.

Some examples of λ -terms:

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

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

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

We observe that there are no numbers or arithmetic operators in the λ -calculus. The calculus can be used to denote and evaluate numerical functions, but the numbers will be represented by λ -terms and the arithmetic operators will be defined by λ -terms.

When we state and prove theorems about λ -calculus we need names standing for arbitrary terms and identifiers. We will use capitals in the middle of the alphabet, sometimes with an index, to range over λ -terms and small Greek letters, possibly indexed, to range over identifiers. Thus there are three kinds of λ -terms:

An *identifier* σ ,

an *abstraction* $(\lambda \alpha. M)$,

and an *application* $(M_1 M_2)$,

A Haskell data type for representing λ -terms is

```
data Lambda = Id String | Abstr String Lambda | Appl Lambda Lambda
```

In Java we would define one abstract class `Lambda` with three subclasses.

Nielson would give the following abstract grammar

$$M ::= x \mid \lambda x. M \mid M_1 M_2$$

Simplified syntax

A complex λ -term will have a lot of parentheses. Parentheses is not very reader friendly. One usually introduces a relaxed syntax for λ -terms where parentheses may be omitted using two associativity rules.

A sequence of applications is assumed to associates to the left so that

$$M_1 M_2 M_3 \dots M_n \triangleq (\dots((M_1 M_2) M_3) \dots M_n)$$

We use the symbol \triangleq when we make definitions.

The abstraction dot, on the other hand, associates to the right so that

$$\lambda\sigma_1.\lambda\sigma_2.\dots\lambda\sigma_n.M \triangleq (\lambda\sigma_1.(\lambda\sigma_2.(\dots(\lambda\sigma_n.M)\dots)))$$

These conventions is expressed in the following concrete grammar.

$$\begin{aligned} \text{term} &\rightarrow \text{factor factor}^* \\ \text{factor} &\rightarrow \text{id} \mid \lambda \text{id} . \text{term} \mid (\text{term}) \end{aligned}$$

The grammar is ambiguous; there are, for example, two derivation trees for $\lambda x . x x$, corresponding to $\lambda x . (x x)$ or $(\lambda x . x) x$. The first one should be chosen; the *term* in the abstraction should extend as far as possible.

An even shorter notation with the same meaning will be used

$$\lambda\sigma_1 \sigma_2 \dots \sigma_n . M \triangleq (\lambda\sigma_1.(\lambda\sigma_2.(\dots(\lambda\sigma_n . M)\dots)))$$

Free and bound identifiers

In $f(x) = x + y$ the identifier x is a formal parameter while y is a identifier whose value is defined elsewhere. In λ -calculus we will call x a *bound* identifier and y a *free* identifier. We define the set of free identifiers in M inductively.

$$\begin{aligned} \mathcal{F}(\sigma) &\triangleq \{\sigma\} \\ \mathcal{F}(\lambda\sigma . M) &\triangleq \mathcal{F}(M) \setminus \{\sigma\} \\ \mathcal{F}(MN) &\triangleq \mathcal{F}(M) \cup \mathcal{F}(N) \end{aligned}$$

The set of bound identifiers is defined likewise.

$$\begin{aligned} \mathcal{B}(\sigma) &\triangleq \emptyset \\ \mathcal{B}(\lambda\sigma . M) &\triangleq \mathcal{B}(M) \cup \{\sigma\} \\ \mathcal{B}(MN) &\triangleq \mathcal{B}(M) \cup \mathcal{B}(N) \end{aligned}$$

A identifier may occur both free an bound in a λ -term, but each occurrence of an identifier is either free or bound. In $(\lambda x . x) x$ the first two occurrences of x is bound while the last is free. We define the set of all identifiers occurring in a λ -term;

$$\mathcal{I}(M) \triangleq \mathcal{F}(M) \cup \mathcal{B}(M)$$

A λ -term without free identifiers is called a *combinator*. Some important combinators have standard names:

$$\begin{aligned} I &\triangleq \lambda x . x \\ K &\triangleq \lambda x . \lambda y . x \\ K_* &\triangleq \lambda x . \lambda y . y \\ S &\triangleq \lambda x . \lambda y . \lambda z . xz(yz) \end{aligned}$$

Conversions

There are two kinds of conversions that may be performed on λ -terms, α - and β -conversions. α -conversion means renaming all bound occurrences of an identifier in an abstraction. β -conversion corresponds to applying a function to an argument and represents one step in the evaluation of a function. The conversions are essentially performed by substitutions. There are however some complications with name collisions.

The basic operation is *substitution*. First we define substitution in λ -terms. We use $M_1[M/\sigma]$ to denote the result after substituting the term M for the identifier σ in the term M_1 . We define it over the structure of M .

$$\begin{aligned} \sigma[M/\sigma] &\triangleq M \\ \tau[M/\sigma] &\triangleq \tau \\ &\text{if } \tau \text{ and } \sigma \text{ are different identifiers} \\ (M_1M_2)[M/\sigma] &\triangleq (M_1[M/\sigma])(M_2[M/\sigma]) \end{aligned}$$

Substitution in abstractions are more complicated since with a naive substitution a free occurrence of an identifier in N may become bound in the abstraction as in $(\lambda x . y)[x/y]$.

$$\begin{aligned} (\lambda \sigma . M)[N/\sigma] &\triangleq (\lambda \sigma . M) \\ (\lambda \tau . M)[N/\sigma] &\triangleq (\lambda \tau . (M[N/\sigma])) \\ &\text{if } \tau \neq \sigma \wedge \tau \notin \mathcal{F}(N) \\ (\lambda \tau . M)[N/\sigma] &\triangleq (\lambda v . ((M[v/\tau])[N/\sigma])) \\ &\text{if } \tau \neq \sigma \wedge \tau \in \mathcal{F}(N) \wedge v \notin \mathcal{I}(M) \cup \mathcal{I}(N) \cup \{\sigma\} \end{aligned}$$

In the last line v is a new identifier not occurring in M or N .

We can formally define the conversions with axioms and inference rules. We start with α -conversion.

$$\lambda \sigma . M \longrightarrow_{\alpha} \lambda \tau . M[\tau/\sigma] \text{ provided that } \tau \notin \mathcal{F}(M)$$

$$\frac{M \longrightarrow_{\alpha} M'}{MN \longrightarrow_{\alpha} M'N} \quad \frac{N \longrightarrow_{\alpha} N'}{MN \longrightarrow_{\alpha} MN'}$$

$$\frac{M \longrightarrow_{\alpha} M'}{\lambda \sigma . M \longrightarrow_{\alpha} \lambda \sigma . M'}$$

We write $M =_{\alpha} N$ if there is a sequence $M = M_0 \longrightarrow_{\alpha} M_1 \longrightarrow_{\alpha} \dots \longrightarrow_{\alpha} M_k = N$ for some $k \geq 0$. $=_{\alpha}$ is an *equivalence relation*, i.e.

1. $M =_\alpha M$, reflexivity
2. if $M =_\alpha N$ then $N =_\alpha M$, symmetry
3. if $M =_\alpha N$ and $N =_\alpha L$ then $M =_\alpha L$, transitivity

We define \rightarrow_β similarly

$$(\lambda\sigma . M)N \rightarrow_\beta M[N/\sigma]$$

$$\frac{M \rightarrow_\beta M'}{MN \rightarrow_\beta M'N} \quad \frac{N \rightarrow_\beta N'}{MN \rightarrow_\beta MN'}$$

$$\frac{M \rightarrow_\beta M'}{\lambda\sigma . M \rightarrow_\beta \lambda\sigma . M'}$$

We write $M \rightarrow_\beta N$ if there is a sequence $M \equiv M_0 \rightarrow_\beta M_1 \rightarrow_\beta \dots \rightarrow_\beta M_k \equiv N$ for some $k \geq 0$. The relation \rightarrow_β is not an equivalence relation, why?

We write $M =_\beta N$ if $M \rightarrow_\beta N$ or $N \rightarrow_\beta M$. The relation $=_\beta$ is an equivalence relation.

Normal form

Let M be a λ -term. If there is no λ -term N such that $M \rightarrow_\beta N$ we say that M is a *normal form*. If a λ -term is not a normal form then it must have a *sub term* that has the form $((\lambda\alpha . M)N)$. Such a sub term will be called a *redex*.

The term $\lambda x . (xy)$ is a normal form and $((\lambda x . x)y)$ may be reduced to y which is a normal form. We say that a λ -term, M , can be *reduced to normal form* if there is a normal form N such that $M \rightarrow_\beta N$.

Some terms cannot be reduced to normal form:

$$(\lambda x . (xx))(\lambda x . (xx)) \rightarrow_\beta (xx)[(\lambda x . (xx))/x] = (\lambda x . (xx))(\lambda x . (xx))$$

Reductions may even make a term more complicated:

$$(\lambda x . (xxx))(\lambda x . (xxx)) \rightarrow_\beta (\lambda x . (xxx))(\lambda x . (xxx))(\lambda x . (xxx))$$

Further reductions will make it even worse.

A term may often be reduced in several ways. If the middle application in

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

is reduced repeatedly we will never reach a normal form while reduction from the left will produce z after two reductions.

The normal form would not deserve its name if a term could be reduced to two essentially different normal forms. Church and Rosser have proved that this is not the case.

Theorem [Church-Rosser]. If $M \rightarrow_\beta N$ and $M \rightarrow_\beta N'$ where N and N' are normal forms then $N =_\alpha N'$. ■

The original proof required several pages. A recent proof by Per Martin-Lf is substantially shorter. The proof consists of several proofs by structural induction. We will omit the proof.

A stronger version of the theorem can be used when the terms cannot be reduced to normal form.

Theorem [Diamond property]. If $M \rightarrow_{\beta} M_0$ and $M \rightarrow_{\beta} M_1$ then there are terms, $M'_0 =_{\alpha} M'_1$, such that $M_0 \rightarrow_{\beta} M'_0$ and $M_1 \rightarrow_{\beta} M'_1$. ■

Henceforth we will write $M = N$ if and only if there are λ -terms, $M' =_{\alpha} N'$, such that $M \rightarrow_{\beta} M'$ and $N \rightarrow_{\beta} N'$.

Two reduction orders are of special interest: *normal order* and *applicative order*. Normal order means that in each step the leftmost redex is reduced. An example:

$$\begin{aligned} (\lambda x . xx)((\lambda y . y)(\lambda z . z)) &\rightarrow_{\beta} ((\lambda y . y)(\lambda z . z))((\lambda y . y)(\lambda z . z)) &\rightarrow_{\beta} \\ (\lambda z . z)((\lambda y . y)(\lambda z . z)) &\rightarrow_{\beta} (\lambda y . y)(\lambda z . z) &\rightarrow_{\beta} \lambda z . z \end{aligned}$$

Using applicative order we compute the argument in an application before reducing the application. The argument may contain applications with arguments which have to be reduced recursively. An example

$$(\lambda x . xx)((\lambda y . y)(\lambda z . z)) \rightarrow_{\beta} (\lambda x . xx)(\lambda z . z) \rightarrow_{\beta} (\lambda z . z)(\lambda z . z) \rightarrow_{\beta} \lambda z . z$$

Church and Rosser also showed that if a term may be reduced to normal form this can be done by using normal order reductions. We have seen that applicative order may fail to terminate when normal order reduction succeeds.

Fixed points

Given a mapping $f \in R \rightarrow R$ we say that x is a *fixed point* of f if f maps x on itself. The fixed points are solutions to the equation $f(x) = x$. The function $f(x) = x^2 - 2$ has two fixed points, -1 and 2 . In same way we say that a λ -term is a fixed point of a combinator F if $FX = X$. The following theorem shows that any combinator has a fixed point and provides a method to find it.

Theorem [Fixed point]. For any combinator F there is a combinator X such that $FX = X$. Furthermore, there is a combinator, $Y \triangleq \lambda f . (\lambda x . f(xx))(\lambda x . f(xx))$, such that YF is a fixed point of F , i.e. $YF = F(YF)$. ■

Proof. Let $W \triangleq \lambda x . F(xx)$ and $X \triangleq WW$. Then

$$X = WW = (\lambda x . F(xx))W = F(WW) = FX$$

Further

$$YF = (\lambda x . F(xx))(\lambda x . F(xx)) = WW = X$$

□

The *fixed point combinator* Y can be used to mimic evaluation of functions defined by recursion in the λ -calculus. We present an informal example using λ -notation with natural numbers, arithmetic operators and conditional expressions. The factorial function is defined by

$$fac(n) \triangleq \text{if } n = 0 \text{ then } 1 \text{ else } n * fac(n - 1);$$

Using λ -notation we can write

$$\begin{aligned} fac &= \lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } n * fac(n - 1) \\ &= (\lambda f. (\lambda n. \text{if } n = 0 \text{ then } 1 \text{ else } n * f(n - 1))) fac = F fac \end{aligned}$$

where

$$F \triangleq \lambda f n. (\text{if } n = 0 \text{ then } 1 \text{ else } n * f(n - 1))$$

Now fac is a fixed point F . Let us investigate if YF is a fixed that represents the factorial function. Let us check that $fac\ 3$ can be reduced to 6.

$$\begin{aligned} fac\ 3 &= (YF)3 = F(YF)3 = \text{if } 3 = 0 \text{ then } 1 \text{ else } 3 * ((YF)2) = 3 * ((YF)2) \\ (YF)2 &= F(YF)2 = \text{if } 2 = 0 \text{ then } 1 \text{ else } 2 * ((YF)1) = 2 * ((YF)1) \\ (YF)1 &= F(YF)1 = \text{if } 1 = 0 \text{ then } 1 \text{ else } 1 * ((YF)0) = 1 * ((YF)0) \\ (YF)0 &= F(YF)0 = \text{if } 0 = 0 \text{ then } 1 \text{ else } 0 * ((YF)0) = 1 \end{aligned}$$

In the following sections we shall represent natural numbers by λ -terms and define combinators that can be used to add and multiply numbers represented in this way. This will make this example fully formal.

Booleans and conditional expressions

Boolean values will be used in conditional expression to choose between alternative expressions. It is convenient to use representations which make this choice simple. Thus we define

$$\begin{aligned} T &\triangleq K = \lambda x. \lambda y. x \\ F &\triangleq K_* = \lambda x. \lambda y. y \end{aligned}$$

A conditional expression $\text{if } B \text{ then } P \text{ else } Q$ may now be represented by BPQ where B is a term that either reduces to T or F . We get

$$\begin{aligned} T\ PQ &= K\ PQ = (\lambda x. \lambda y. x)PQ = P \\ F\ PQ &= K_*\ PQ = (\lambda x. \lambda y. y)PQ = Q \end{aligned}$$

The logical terms T and F can be applied to two terms and select the first or the second.

Pairs and natural numbers

We can represent an ordered pair as a λ -term:

$$[M, N] \triangleq \lambda z. zMN$$

and extract the components with T and F .

$$\begin{aligned} [M, N]T &= M \\ [M, N]F &= N \end{aligned}$$

We use pairs to represent natural numbers in a way that makes the successor function simple. We will use $\lceil n \rceil$ as a name of the λ -term which represents the natural number n and define $\lceil n \rceil$ by induction:

$$\begin{aligned} \lceil 0 \rceil &\triangleq I = \lambda x. x \\ \lceil n + 1 \rceil &\triangleq [F, \lceil n \rceil] \end{aligned}$$

There are simple combinators that represents the successor function $suc\ n = n + 1$, the predecessor function $pred\ n = n - 1$ and the predicate $iszero\ n = (n = 0)$:

$$\begin{aligned} Suc &\triangleq \lambda x. [F, x] \\ Pred &\triangleq \lambda x. x F \\ IsZero &\triangleq \lambda x. x T \end{aligned}$$

We leave it as an exercise to show that these terms are adequate representations.

We proceed as we did for the factorial function to get a representation for addition.

$$add\ n\ m = \text{if } m = 0 \text{ then } n \text{ else } suc(add\ n\ (pred\ m));$$

Translation to λ -calculus

$$Add \triangleq \lambda n. \lambda m. IsZero\ m\ n\ (Suc(Add\ n\ (Pred\ m)))$$

This is rewritten as a fixed point equation

$$Add = F\ Add$$

where

$$F = \lambda f. \lambda n. \lambda m. IsZero\ m\ n\ (Suc(f\ n\ (Pred\ m)))$$

with the solution $Add = YF$.

In a similar way we may define a combinator Mul such that $Mul\ \lceil m \rceil\ \lceil n \rceil = \lceil m * n \rceil$. With these in hand a factorial combinator may be constructed.

