

A review of the simply typed λ -calculus

Andrej Bauer

March 24, 2022

The λ -calculus is the abstract theory of functions, just like group theory is the abstract theory of symmetries. There are two basic operations that can be performed with functions. The first one is the *application* of a function to an argument: if f is a function and a is an argument, then fa is the application of f to a . The second operation is *abstraction*: if x is a variable and t is an expression in which x may appear, then there is a function f defined by

$$fx = t .$$

Here we gave the name f to the newly formed function. But we could have expressed the same function without giving it a name; this is usually written as

$$x \mapsto t ,$$

and it means “ x is mapped to t ”. In λ -calculus we use a different notation, which is more convenient when abstractions are nested:

$$\lambda x . t .$$

This operation is called λ -*abstraction*. For example, $\lambda x . \lambda y . (x + y)$ is the function which maps an argument a to the function $\lambda y . (a + y)$.

In an expression $\lambda x . t$ the variable x is *bound* in t .

There are two kinds of λ -calculus, the *typed* and the *untyped* one. In the untyped version there are no restrictions on how application is formed, so that an expression such as

$$\lambda x . (xx)$$

is valid, whatever it means. In typed λ -calculus every expression has a *type*, and there are rules for forming valid expressions and types. For example, we can only form an application f, a when a has a type A and f has a type $A \rightarrow B$, which indicates a function taking arguments of type A and giving

results of type B . The judgment that expression t has a type A is written as

$$t : A .$$

To computer scientists the idea of expressions having types is familiar from programming languages, whereas mathematicians can think of types as sets and read $t : A$ as $t \in A$. We will concentrate on the typed λ -calculus.

We now give a precise definition of what constitutes a *simply-typed λ -calculus*. First, we are given a set of *simple types*, which are generated from *basic types* by formation of products and function types:

$$\begin{aligned} \text{Basic type } B &::= B_0 \mid B_1 \mid B_2 \cdots \\ \text{Simple type } A &::= B \mid A_1 \times A_2 \mid A_1 \rightarrow A_2. \end{aligned}$$

Function types associate to the right:

$$A \rightarrow B \rightarrow C \equiv A \rightarrow (B \rightarrow C) .$$

We assume there is a countable set of *variables* x, y, u, \dots . We are also given a set of *basic constants*. The set of *terms* is generated from variables and basic constants by the following grammar:

$$\begin{aligned} \text{Variable } v &::= x \mid y \mid z \mid \cdots \\ \text{Constant } c &::= c_1 \mid c_2 \mid \cdots \\ \text{Term } t &::= v \mid c \mid * \mid \langle t_1, t_2 \rangle \mid \mathbf{fst} \, t \mid \mathbf{snd} \, t \mid t_1 \, t_2 \mid \lambda x : A . t \end{aligned}$$

In words, this means:

1. a variable is a term,
2. each basic constant is a term,
3. the constant $*$ is a term, called the *unit*,
4. if u and t are terms then $\langle u, t \rangle$ is a term, called a *pair*,
5. if t is a term then $\mathbf{fst} \, t$ and $\mathbf{snd} \, t$ are terms,
6. if u and t are terms then $u \, t$ is a term, called an *application*
7. if x is a variable, A is a type, and t is a term, then $\lambda x : A . t$ is a term, called a *λ -abstraction*.

The variable x is *bound* in $\lambda x : A . t$. Application associates to the left, thus $st u = (st)u$. The *free variables* $\text{FV}(t)$ of a term t are computed as follows:

$$\begin{aligned} \text{FV}(x) &= \{x\} && \text{if } x \text{ is a variable} \\ \text{FV}(a) &= \emptyset && \text{if } a \text{ is a basic constant} \\ \text{FV}(\langle u, t \rangle) &= \text{FV}(u) \cup \text{FV}(t) \\ \text{FV}(\mathbf{fst} t) &= \text{FV}(t) \\ \text{FV}(\mathbf{snd} t) &= \text{FV}(t) \\ \text{FV}(u t) &= \text{FV}(u) \cup \text{FV}(t) \\ \text{FV}(\lambda x . t) &= \text{FV}(t) \setminus \{x\} . \end{aligned}$$

If x_1, \dots, x_n are *distinct* variables and A_1, \dots, A_n are types then the sequence

$$x_1 : A_1, \dots, x_n : A_n$$

is a *typing context*, or just *context*. The empty sequence is sometimes denoted by a dot \cdot , and it is a valid context. Context are denoted by capital Greek letters Γ, Δ, \dots

A *typing judgment* is a judgment of the form

$$\Gamma \mid t : A$$

where Γ is a context, t is a term, and A is a type. In addition the free variables of t must occur in Γ , but Γ may contain other variables as well. We read the above judgment as “in context Γ the term t has type A ”. Next we describe the rules for deriving typing judgments.

Each basic constant c_i has a uniquely determined type C_i ,

$$\overline{\Gamma \mid c_i : C_i}$$

The type of a variable is determined by the context:

$$\overline{x_1 : A_1, \dots, x_i : A_i, \dots, x_n : A_n \mid x_i : A_i} \quad (1 \leq i \leq n)$$

The constant $*$ has type 1 :

$$\overline{\Gamma \mid * : 1}$$

The typing rules for pairs and projections are:

$$\frac{\Gamma \mid u : A \quad \Gamma \mid t : B}{\Gamma \mid \langle u, t \rangle : A \times B} \quad \frac{\Gamma \mid t : A \times B}{\Gamma \mid \mathbf{fst} t : A} \quad \frac{\Gamma \mid t : A \times B}{\Gamma \mid \mathbf{snd} t : B}$$

The typing rules for application and λ -abstraction are:

$$\frac{\Gamma \mid t : A \rightarrow B \quad \Gamma \mid u : A}{\Gamma \mid tu : B} \qquad \frac{\Gamma, x : A \mid t : B}{\Gamma \mid (\lambda x : A. t) : A \rightarrow B}$$

Lastly, we have *equations* between terms; for terms of type A in context Γ ,

$$\Gamma \mid u : A, \qquad \Gamma \mid t : B,$$

the judgment that they are equal is written as

$$\Gamma \mid u = t : A.$$

Note that u and t necessarily have the same type; it does *not* make sense to compare terms of different types. We have the following rules for equations:

1. Equality is an equivalence relation:

$$\frac{}{\Gamma \mid t = t : A} \qquad \frac{\Gamma \mid t = u : A}{\Gamma \mid u = t : A} \qquad \frac{\Gamma \mid t = u : A \quad \Gamma \mid u = v : A}{\Gamma \mid t = v : A}$$

2. The weakening rule:

$$\frac{\Gamma \mid u = t : A}{\Gamma, x : B \mid u = t : A}$$

3. Unit type:

$$\frac{}{\Gamma \mid t = * : 1}$$

4. Equations for product types:

$$\frac{\Gamma \mid u = v : A \quad \Gamma \mid s = t : B}{\Gamma \mid \langle u, s \rangle = \langle v, t \rangle : A \times B}$$

$$\frac{\Gamma \mid s = t : A \times B}{\Gamma \mid \mathbf{fst} s = \mathbf{fst} t : A} \qquad \frac{\Gamma \mid s = t : A \times B}{\Gamma \mid \mathbf{snd} s = \mathbf{snd} t : A}$$

$$\frac{}{\Gamma \mid t = \langle \mathbf{fst} t, \mathbf{snd} t \rangle : A \times B}$$

$$\frac{}{\Gamma \mid \mathbf{fst} \langle u, t \rangle = u : A} \qquad \frac{}{\Gamma \mid \mathbf{snd} \langle u, t \rangle = t : A}$$

5. Equations for function types:

$$\frac{\Gamma \mid s = t : A \rightarrow B \quad \Gamma \mid u = v : A}{\Gamma \mid su = tv : B}$$

$$\frac{\Gamma, x : A \mid t = u : B}{\Gamma \mid (\lambda x : A. t) = (\lambda x : A. u) : A \rightarrow B}$$

$$\frac{}{\Gamma \mid (\lambda x : A. t)u = t[u/x] : A} \quad (\beta\text{-rule})$$

$$\frac{}{\Gamma \mid \lambda x : A. (tx) = t : A \rightarrow B} \quad \text{if } x \notin \mathbf{FV}(t) \quad (\eta\text{-rule})$$

This completes the description of a simply-typed λ -calculus.

Apart from the above rules for equality we might want to impose additional equations. In this case we do not speak of a λ -calculus but rather of a λ -theory. Thus, a λ -theory \mathbb{T} is given by a set of basic types, a set of basic constants, and a set of *equations* of the form

$$\Gamma \mid u = t : A .$$

We summarize the preceding definitions.

Definition 1 A *simply-typed λ -calculus* is given by a set of *basic types* and a set of *basic constants* together with their types. A *simply-typed λ -theory* is a simply-typed λ -calculus together with a set of equations.

We use letters $\mathbb{S}, \mathbb{T}, \mathbb{U}, \dots$ to denote theories.

Example 2 The theory of a group is a simply-typed λ -theory. It has one basic type \mathbf{G} and three basic constant, the unit \mathbf{e} , the inverse \mathbf{i} , and the group operation \mathbf{m} ,

$$\mathbf{e} : \mathbf{G} , \quad \mathbf{i} : \mathbf{G} \rightarrow \mathbf{G} , \quad \mathbf{m} : \mathbf{G} \times \mathbf{G} \rightarrow \mathbf{G} ,$$

with the following equations:

$$\begin{aligned} x : \mathbf{G} \mid \mathbf{m}\langle x, \mathbf{e} \rangle &= x : \mathbf{G} \\ x : \mathbf{G} \mid \mathbf{m}\langle \mathbf{e}, x \rangle &= x : \mathbf{G} \\ x : \mathbf{G} \mid \mathbf{m}\langle x, \mathbf{i} x \rangle &= \mathbf{e} : \mathbf{G} \\ x : \mathbf{G} \mid \mathbf{m}\langle \mathbf{i} x, x \rangle &= \mathbf{e} : \mathbf{G} \\ x : \mathbf{G}, y : \mathbf{G}, z : \mathbf{G} \mid \mathbf{m}\langle x, \mathbf{m}\langle y, z \rangle \rangle &= \mathbf{m}\langle \mathbf{m}\langle x, y \rangle, z \rangle : \mathbf{G} \end{aligned}$$

These are just the familiar axioms for a group.

Example 3 In general, any algebraic theory \mathbb{A} determines a λ -theory. There is one basic type \mathbf{A} and for each operation f of arity k there is a basic constant $\mathbf{f} : \mathbf{A}^k \rightarrow \mathbf{A}$, where \mathbf{A}^k is the k -fold product $\mathbf{A} \times \dots \times \mathbf{A}$. It is understood that $\mathbf{A}^0 = 1$. The terms of \mathbb{A} are translated to the terms of the corresponding λ -theory in a straightforward manner. For every axiom $t = u$ of \mathbb{A} the corresponding axiom in the λ -theory is

$$x_1 : \mathbf{A}, \dots, x_n : \mathbf{A} \mid t = u : \mathbf{A}$$

where x_1, \dots, x_n are the variables occurring in t and u .

Example 4 The theory of a directed graph is a simply-typed theory with two basic types, V for vertices and E for edges, and two basic constant, source src and target trg ,

$$\text{src} : E \rightarrow V, \quad \text{trg} : E \rightarrow V.$$

There are no equations.

Example 5 An example of a λ -theory is readily found in the theory of programming languages. The mini-programming language PCF is a simply-typed λ -calculus with a basic type nat for natural numbers, and a basic type bool of Boolean values,

$$\text{Basic type } B ::= \text{nat} \mid \text{bool}.$$

There are basic constants zero 0 , successor succ , the Boolean constants true and false , comparison with zero iszero , and for each type A the *conditional* cond_A and the *fixpoint* operator fix_A . They have the following types:

$$\begin{aligned} 0 &: \text{nat} \\ \text{succ} &: \text{nat} \rightarrow \text{nat} \\ \text{true} &: \text{bool} \\ \text{false} &: \text{bool} \\ \text{iszero} &: \text{nat} \rightarrow \text{bool} \\ \text{cond}_A &: \text{bool} \rightarrow A \rightarrow A \\ \text{fix}_A &: (A \rightarrow A) \rightarrow A \end{aligned}$$

The equational axioms of PCF are:

$$\begin{aligned} &\cdot \mid \text{iszero } 0 = \text{true} : \text{bool} \\ x : \text{nat} &\mid \text{iszero } (\text{succ } x) = \text{false} : \text{bool} \\ u : A, t : A &\mid \text{cond}_A \text{ true } u t = u : A \\ u : A, t : A &\mid \text{cond}_A \text{ false } u t = t : A \\ t : A \rightarrow A &\mid \text{fix}_A t = t(\text{fix}_A t) : A \end{aligned}$$

Example 6 Another example of a λ -theory is the *theory of a reflexive type*. This theory has one basic type D and two constants

$$\text{r} : D \rightarrow D \rightarrow D \quad \text{s} : (D \rightarrow D) \rightarrow D$$

satisfying the equation

$$f : D \rightarrow D \mid \mathbf{r}(\mathbf{s} f) = f : D \rightarrow D \quad (1)$$

which says that \mathbf{s} is a section and \mathbf{r} is a retraction, so that the function type $D \rightarrow D$ is a subspace (even a retract) of D . A type with this property is said to be *reflexive*. We may additionally stipulate the axiom

$$x : D \mid \mathbf{s}(\mathbf{r} x) = x : D \quad (2)$$

which implies that D is isomorphic to $D \rightarrow D$.

Untyped λ -calculus We briefly describe the *untyped λ -calculus*. It is a theory whose terms are generated by the following grammar:

$$t ::= v \mid t_1 t_2 \mid \lambda x. t .$$

In words, a variable is a term, an application $t t'$ is a term, for any terms t and t' , and a λ -abstraction $\lambda x. t$ is a term, for any term t . Variable x is bound in $\lambda x. t$. A *context* is a list of distinct variables,

$$x_1, \dots, x_n .$$

We say that a term t is valid in context Γ if the free variables of t are listed in Γ . The judgment that two terms u and t are equal is written as

$$\Gamma \mid u = t ,$$

where it is assumed that u and t are both valid in Γ . The context Γ is not really necessary but we include it because it is always good practice to list the free variables.

The rules of equality are as follows:

1. Equality is an equivalence relation:

$$\frac{}{\Gamma \mid t = t} \quad \frac{\Gamma \mid t = u}{\Gamma \mid u = t} \quad \frac{\Gamma \mid t = u \quad \Gamma \mid u = v}{\Gamma \mid t = v}$$

2. The weakening rule:

$$\frac{\Gamma \mid u = t}{\Gamma, x \mid u = t}$$

3. Equations for application and λ -abstraction:

$$\frac{\Gamma \mid s = t \quad \Gamma \mid u = v}{\Gamma \mid s u = t v} \quad \frac{\Gamma, x \mid t = u}{\Gamma \mid \lambda x. t = \lambda x. u} \quad (\beta\text{-rule})$$

$$\frac{}{\Gamma \mid (\lambda x. t)u = t[u/x]} \quad (\beta\text{-rule})$$

$$\frac{}{\Gamma \mid \lambda x. (t x) = t} \quad \text{if } x \notin \text{FV}(t) \quad (\eta\text{-rule})$$

The untyped λ -calculus can be translated into the theory of a reflexive type from Example 6. An untyped context Γ is translated to a typed context Γ^* by typing each variable in Γ with the reflexive type D , i.e., a context x_1, \dots, x_k is translated to $x_1 : \mathsf{D}, \dots, x_k : \mathsf{D}$. An untyped term t is translated to a typed term t^* as follows:

$$\begin{aligned} x^* &= x && \text{if } x \text{ is a variable,} \\ (ut)^* &= (\mathbf{r} u^*)t^*, \\ (\lambda x. t)^* &= \mathbf{s}(\lambda x : \mathsf{D}. t^*). \end{aligned}$$

For example, the term $\lambda x. (x x)$ translates to $\mathbf{s}(\lambda x : \mathsf{D}. ((\mathbf{r} x) x))$. A judgment

$$\Gamma \mid u = t \quad (3)$$

is translated to the judgment

$$\Gamma^* \mid u^* = t^* : \mathsf{D}. \quad (4)$$

Exercise* 7 Prove that if equation (3) is provable then equation (4) is provable as well. Identify precisely at which point in your proof you need to use equations (1) and (2). Does provability of (4) imply provability of (3)?