

Type theory

Simon Charlow (simon.charlow@rutgers.edu)

September 30, 2015

1 Lambda calculus, redux

Why we bother:

- λ s let us build functions **anonymously**.
- Simplification of λ -terms is a straightforward way to compute values for complex expressions.

1.1 Syntax

Syntax: the set of λ -terms, Λ , is the smallest set meeting the following conditions:¹

1. (**Variables**) If x is a variable, $x \in \Lambda$.
2. (**Abstraction**) If x is a variable and $M \in \Lambda$, then $(\lambda x. M) \in \Lambda$.
3. (**Application**) If $M, N \in \Lambda$, then $(M N) \in \Lambda$.

And that is literally it. But some abbreviatory conventions improve usability:

1. Application is left-associative. For example:
 - ‘ $M N L$ ’ instead of ‘ $((M N) L)$ ’
2. Parentheses can be omitted whenever doing so doesn’t create ambiguity. For example:
 - ‘ $\lambda x. M N$ ’ instead of ‘ $\lambda x. (M N)$ ’
 - ‘ $\lambda x. M$ ’ instead of ‘ $(\lambda x. M)$ ’
 - ‘ $\lambda x. \lambda y. M$ ’ instead of ‘ $(\lambda x. (\lambda y. M))$ ’

NB: $\lambda x. M y$ is equivalent to $\lambda x. (M y)$ and $(\lambda x. M) y$, but *not* $(\lambda x. M) y$. I.e., the scope of a λ always extends as far to the right as possible (i.e. while still yielding a well-formed term).

3. A sequence of abstractions can be abbreviated. For example:
 - ‘ $\lambda x y z. M$ ’ instead of ‘ $\lambda x. \lambda y. \lambda z. M$ ’

I recommend mastering at least the first two conventions. Whether you use the third is up to you...

¹As we saw in the last class, in BNF this amounts to $E ::= V \mid (\lambda V. E) \mid (E_1 E_2)$.

1.2 “Semantics”

Core equivalences:

- **(Renaming bound variables)** $\lambda x. M \xrightarrow{\alpha} \lambda y. M[x \rightarrow y]$
- **(Applying functions to arguments)** $(\lambda x. M) y \xrightarrow{\beta} M[x \rightarrow y]$
- **(Extensionality)** $f \xrightarrow{\eta} \lambda x. f x$

Two of these equivalences rely on the notion of **substitution**, defined as follows:

$M[x \rightarrow y]$ is the formula you get by replacing all M 's **free** occurrences of x with y .

The **free** variables in an expression E ('FVE') are the variables in E not **bound** by a λ :

$$\begin{aligned} \text{FV } x &= \{x\} \\ \text{FV } (\lambda x. M) &= \text{FV } M - \{x\} \\ \text{FV } (M N) &= \text{FV } M \cup \text{FV } N \end{aligned}$$

In other words, you can read $M[x \rightarrow y]$ as “ M , but where the (free) x 's becomes y 's”.

Some examples of substitution in action:

$$\begin{aligned} x[x \rightarrow y] &= y \\ (\lambda x. x)[x \rightarrow y] &= \lambda x. x \\ (\lambda f. f x)[x \rightarrow y] &= \lambda f. f y \end{aligned}$$

Substitution has various notations. Sometimes you see ' $M[x := y]$ '. Sometimes, as in our last class, you see ' $M[y / x]$ '. I'm hoping the present choice will be the most perspicuous. Let me know.

A couple examples of α , β , and η in action (notice that α -, β -, and η - equivalence can apply to proper *sub-formulas* of any λ -expression, as well as maximal λ -expressions):

$$\begin{aligned} \lambda f. \lambda x. f x &\xrightarrow{\alpha} \lambda f. \lambda z. f z \xrightarrow{\alpha} \lambda g. \lambda z. g z \\ (\lambda f. \lambda x. f x) (\lambda y. y) &\xrightarrow{\beta} \lambda x. (\lambda y. y) x \xrightarrow{\beta} \lambda x. x \\ \lambda f. f &\xrightarrow{\eta} \lambda f. \lambda x. f x \xrightarrow{\eta} \lambda f. \lambda x. \lambda y. f x y \end{aligned}$$

A β -reduction is **invalid** if it places a variable within the scope of a co-indexed λ :

$$(\lambda x. \lambda y. f x y) y \xrightarrow{\beta? \text{NO.}} \lambda y. f y y$$

To avoid this unfortunate situation, exploit α -equivalence:

$$(\lambda x. \lambda y. f x y) y \xrightarrow{\alpha} (\lambda x. \lambda z. f x z) y \xrightarrow{\alpha} \lambda z. f y z$$

An α -equivalence is **invalid** for an expression E if it binds a free variable in E :

$$\lambda x. x y \xrightarrow{\alpha? \text{NO.}} \lambda y. y y$$

To avoid this unfortunate situation... Don't get into it!

2 Types

The variety of the λ -calculus defined in the last section is **untyped**. Because of this, you can define functions like the following:

$$\omega := \lambda x. x x$$

Which you can use to define the following:

$$\Omega := \omega \omega$$

As we saw in the last class, Ω has an interesting property: β -“reducing” it gives you exactly what you already had!!

$$\Omega \xrightarrow{\beta} \Omega$$

We’ll be working with a different system in this class, called the **simply-typed** λ -calculus.

What are **types**? Pretty simple: types are nothing more than **sets of objects**! For example, the type \mathbb{N} is the set of natural numbers.

In semantics, we work with two basic types: the type of individuals (roughly, things that can go in subject or object position²), and the type of truth values. By convention, the type of individuals is written ‘e’ (for ‘entity’, probably), and the type of truth values is written ‘t’.

$$\begin{aligned} e &::= \{x : x \text{ is an individual}\} \\ t &::= \{0, 1\} \end{aligned}$$

Using our basic types, we can give an inductive definition for the set of **all types** τ , as follows:

1. $e, t \in \tau$.
2. If $\alpha, \beta \in \tau$, then $\langle \alpha, \beta \rangle \in \tau$.
3. Nothing else is in τ .

Some examples of types according to this “grammar”:

- e
- $\langle e, t \rangle$
- $\langle \langle e, t \rangle, t \rangle$
- $\langle \langle e, e \rangle, \langle \langle e, e \rangle, \langle e, e \rangle \rangle \rangle$

The simple types are the individuals and the truth values. What are the complex types? The types of **functions**: $\langle \alpha, \beta \rangle$ is the type of functions from things of type α to things of type β .

An example. Let’s begin with the set of linguists. This set has a characteristic function, which we might name LINGUIST. This function can be represented in λ -notation as follows (notice that this is just the η -expansion of LINGUIST!):

$$\lambda x. \text{LINGUIST } x$$

²NB: semanticists think of lots of things as individuals, including numbers, tables and chairs, pets, bacteria, and so on.

This function takes individuals as arguments and returns truth values. Accordingly, its type is $\langle e, t \rangle$.

Likewise, an $\langle \langle e, t \rangle, t \rangle$ function takes functions from individuals to truth values as arguments, and returns truth values. A trivial example is the following:

$$\lambda p. 1$$

In the simply-typed λ -calculus, every function is **partial**. That is, every function expects its arguments to be of a certain type, and doesn't know what to do otherwise. For example, $\lambda x. \text{LINGUIST } x$ wouldn't know what to do if you fed it $\lambda x. \text{PHILOSOPHER } x$. The result would be **undefined**.

There are a number of ways to notate restrictions on types. Heim & Kratzer use the following convention (more generally, they use the ':' notation to indicate partiality):

$$\lambda x : x \in e. \text{LINGUIST } x$$

Another way to go (the way I prefer, maybe because it saves symbols), is to use subscripts:

$$\lambda x_e. \text{LINGUIST } x$$

In any case, type notations are often omitted when context makes clear what an expression's type is.

Some other examples:

- $\llbracket \text{blue} \rrbracket = \lambda x_e. \text{BLUE } x$
- $\llbracket \text{devoured} \rrbracket = \lambda x_e. \lambda y_e. \text{DEVoured } x y$
- $\llbracket \text{not} \rrbracket = \lambda p_t. 1 - p$
- $\llbracket \text{and} \rrbracket = \lambda p_t. \lambda q_t. p = q = 1$

In order:

- The type of a simple adjective is a function from individuals to truth values, type $\langle e, t \rangle$.
- The type of a transitive verb is a function from individuals to: a function from individuals to truth values, type $\langle e, \langle e, t \rangle \rangle$.
- The type of sentential negation is a function from truth values to truth values, type $\langle t, t \rangle$.
- The type of sentential conjunction is a function from truth values to: a function from truth values to truth values, type $\langle t, \langle t, t \rangle \rangle$.

In this manner of speaking, we can talk about denotations having types, as well as syntactic units (i.e., the type associated with a piece of syntax's denotation). We can use the following notation to indicate that an expression has a certain type:

$$\text{devoured} :: \langle e, \langle e, t \rangle \rangle$$

At the risk of burdening you with another notational convention, I'll mention that the way I prefer to write types in my own work (because it's way easier to parse than the brackets)³ is using arrows:

$$\begin{aligned} \llbracket e \rrbracket &= e \\ \llbracket e \rightarrow t \rrbracket &= e \rightarrow t \\ \llbracket \langle e, \langle e, t \rangle \rangle \rrbracket &= e \rightarrow e \rightarrow t \\ \llbracket \langle \langle e, t \rangle, t \rangle \rrbracket &= (e \rightarrow t) \rightarrow t \end{aligned}$$

³In fact, you could *omit* the arrows entirely and still have a perfectly good notation scheme for types. But that's getting a bit telegraphic for most people's tastes.

Notice that in the simply-typed λ -calculus, it **makes no sense** for a function to apply to itself:

- Take an arbitrary f . Since f is a function, it will have type $\langle \alpha, \beta \rangle$, for some types α and β .
- But then, for $f f$ to be defined, f needs to be able to take itself as an argument.
- So then the functor f 's type must be $\langle \langle \alpha, \beta \rangle, \beta \rangle$.
- But then the argument f 's type must also be $\langle \langle \alpha, \beta \rangle, \beta \rangle$.
- In which case f 's type must actually be $\langle \langle \langle \alpha, \beta \rangle, \beta \rangle, \beta \rangle$.
- And so on...

This means that ω , which crucially involves self-application, is impossible to define in a simply-typed system. Ergo, Ω is impossible to define.

In fact, every simply-typed λ -term is guaranteed to **terminate** into a unique **normal form**, in which no further β -reductions are possible. Because every expression has a *unique* normal form, the order of reductions never makes any difference. For example:

$$\begin{aligned} (\lambda x. f x) ((\lambda y. g y) z) &\xrightarrow{\beta} (\lambda x. f x) (g z) \xrightarrow{\beta} f (g z) \\ (\lambda x. f x) ((\lambda y. g y) z) &\xrightarrow{\beta} f ((\lambda y. g y) z) \xrightarrow{\beta} f (g z) \end{aligned} \tag{1}$$

Unlike ω , it is possible to assign types to all the expressions in this term in a way that achieves a defined end result. Exercise: come up with one such typing scheme.