

## 1 Introduction

What is a program? Is it just something that tells the computer what to do? That misses the point. If a program were just a simple list of instructions we wouldn't be able to do much useful with computers. Programs are much more interesting than that.

And if we think about a program as just a description of what to do, then part of what a program means is baked into the computer that interprets it. So you need to describe that computer too, and *that's* clearly more interesting than just a list of instructions.

Programs describe computation, but they're really a kind of information as well. Suppose we step back from my initial question and ask a truly fundamental question: "what is information?". For most of the information we normally think about, familiar mathematical tools like sets and sequences are adequate to describe that information. Programs describe computation, but programs are information too—a particularly interesting kind of information, especially for computer scientists. In this course we'll see some of the mathematical, formal tools that are needed to precisely describe what programs are, what programming languages are, and maybe better answers to the two questions above.

This is a course about the *semantics* of programs and programming languages. "Semantics" is simply a fancy way of saying "meaning"; we want to give a precise meaning to programs. This is useful for many reasons. It means that we can understand what they are going to do, which is clearly helpful in writing programs that work. It also helps us when building tools like compilers, optimizers, and interpreters because it means we can judge whether these tools are implemented correctly.

There are three major components to this course.

- Dynamic Semantics. We will study methods for describing what a program does when it runs.
- Static Semantics. We will also study methods for reasoning about programs before they run. We would like to rule out programs that are bad in various senses before they run. For example, programs that contain run-time (dynamic) type errors.
- Language features. We will apply methods for dynamic and static semantics to study actual language features of interest, probably including some interesting features that many students have not seen before.

We would like to characterize the semantics of a program as a function produces the output based on inputs. More generally, real programs are reactive and interact with their inputs arriving from the environment. Describing this kind of program is more challenging, though we can view the reactive behavior of the program again as a function of the inputs it receives from the environment as it runs. Thus, to describe program semantics, we will build up some mathematical tools for constructing and reasoning about functions.

### 1.1 Review of mathematical functions

A (*total*) *function* is a relation, in which each element of a set  $A$  is associated with a unique element of another (possibly the same) set  $B$ . If  $f$  is such a function, we write:

$$f : A \rightarrow B$$

Here,  $A$  is the *domain* of the function and  $B$  the *codomain*. The *image* of the function is the set of elements in  $B$  that at least one element in  $A$  is mapped to.

$$\text{image}(f) = \{x \in B : x = f(a) \text{ for some } a \in A\}$$

We avoid the word *range* because it is ambiguous: sometimes it is used to refer to the codomain, sometimes to the image. A *partial function*  $f : A \rightharpoonup B$  is a relation that is a function on some subset of the domain  $\text{dom}(f) \subseteq A$ .

## 1.2 Function representation

One way of describing functions mathematically is via their *extension*: essentially, the set of all pairs of (input, output) values. Other way is to give an *intensional* representation—a program, or what you have to do in order to evaluate the function.

As an aside, this class is quite a bit about how to get from an intensional representation of a program to some more extensional representation of the program, often a function. In this class, program texts actually refer to their parsed abstract syntax trees (AST)

## 1.3 Lambda Calculus

Introduced by Alonzo Church and Stephen Cole Kleene in the 1930s to describe functions in an unambiguous and compact manner, lambda calculus<sup>1</sup> gives an *intensional* representation for a function. The lambda calculus is universal in the sense that any computable function can be expressed and evaluated using this formalism. It is thus equivalent to a *Turing machine*.

Real programming languages such as Lisp, Scheme, Haskell and ML are very much based on lambda calculus, although there are differences as well.

## 2 Untyped Lambda Calculus

### 2.1 Syntax

The following is the syntax of the  $\lambda$ -calculus. A term is defined as follows:

$$\begin{aligned} e ::= & x \quad x \in \mathbf{Var} \quad \text{where } \mathbf{Var} \text{ is some countable set} \\ | & e_0 e_1 \quad (\text{application}) \\ | & \lambda x. e \quad (\text{abstraction}) \end{aligned}$$

In an abstraction,  $x$  is the *argument*,  $e$  is the *body* of the function. Notice that this function doesn't have a name. In mathematics it is common to define functions as  $f(x) = x^2$ . A corresponding anonymous notation that is frequently used is  $x \mapsto x^2$ ; this is really the same thing as  $\lambda x. x^2$ . One nice thing about lambda terms is that they are an *anonymous* representation of functions: the function doesn't have to have a name like  $f$  in order to talk about this.

Here are some examples of terms. First, we present (and define) the identity function:

$$ID = \lambda x. x$$

The next example is a function that will ignore its argument and return the identity function.

$$\lambda x. \lambda a. a$$

This is the same as

$$\lambda x. ID$$

Parentheses are used to show explicitly how to parse expressions, but they are not strictly necessary. It's a good idea to include them if you aren't sure. Lambda terms are like sums: they extend as far to the right as they can. For example,  $(\lambda x. x \lambda y. y)$  is the same thing as  $(\lambda x. (x \lambda y. y))$ , not  $((\lambda x. x) (\lambda y. y))$ . Application is *left-associative* - that is,  $e_1 e_2 e_3$  is the same thing as  $(e_1 e_2) e_3$ .

---

<sup>1</sup>Why  $\lambda$ ? Church wanted to separate the bound variables from the unbound (free) variables and he used a caret on top of the bound variables.  $f(x) = x + yx^2$  was represented as  $\hat{x}.x + yx^2$ . Apparently, the printers could not handle the caret and it moved to the front and became a  $\lambda$ ; the expression became  $\lambda x. x + yx^2$

## 2.2 Closed terms

An occurrence of a variable  $x$  in a term is said to be *bound* if there is an enclosing  $\lambda x. e$ ; otherwise, it is *free*. A *closed term* is one in which all identifiers are bound.

Consider the following term:

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

Both occurrences of  $x$  are bound, the first occurrence of  $y$  is bound, the  $a$  is free, and the last  $y$  is also free, since it is outside the scope of the  $\lambda y$ .

If a program has some variables that are free, then you do not have a complete program as you do not know what to do with the free variables. Hence, a well formed program in lambda calculus is a closed term.

## 2.3 Higher-order functions

In lambda calculus, we can define functions that can take functions as arguments and/or return functions as results. In fact, every argument is a function and every result is a function. That is, functions are first-class values.

This example takes a function as an argument and applies it to 5:

$$\lambda f. (f 5)$$

We can further generalize. This function takes an argument  $v$  and returns a function that calls its argument on  $v$ :

$$\lambda v. \lambda f. (f v)$$

## 2.4 Multi-argument functions and currying

A multi-argument lambda calculus (functions having multiple arguments) could be defined as

$$\begin{aligned} e ::= & \dots \text{earlier stuff}\dots \\ | & \lambda x_1 \dots x_n. e \quad (\text{multi-argument abstraction}) \\ | & e_0 e_1 e_2 \dots e_n \quad (\text{multi-argument application}) \end{aligned}$$

In the multi-argument application,  $e_0$  is a  $n$ -argument function and  $e_1 \dots e_n$  are the arguments.

It is known that this isn't any more expressive than basic lambda calculus. So, we stick to basic lambda calculus but we may occasionally write down terms that look like the above for convenience. It's just *syntactic sugar*. It can be transformed into something more basic that does not require the extended syntax. A transformation that removes syntactic sugar is called *desugaring*.

Desugaring multi-argument functions:

$$\begin{aligned} \lambda x_1 \dots x_n. e & \Rightarrow \lambda x_1. \lambda x_2. \lambda x_n. e \\ e_0 e_1 \dots e_n & \Rightarrow (\dots ((e_0 e_1) e_2) \dots e_n) \end{aligned}$$