

1 Term Equivalence

When are two terms equal? This is not as simple a question as it may seem. As *intensional* objects, two terms are equal if they are syntactically identical. As *extensional* objects, however, two terms should be equal if they represent the same function. We will say that two terms are *equivalent* if they are equal in an extensional sense.

For example, it seems clear that the terms $\lambda x.x$ and $\lambda y.y$ are equivalent. The name of the variable is not essential. But we also probably think that $\lambda x.(\lambda y.y) x$ is equivalent to $\lambda x.x$ too, in a less trivial sense. And there are even more interesting cases, like $\lambda x.\lambda y.x y$.

But what function does a term like $\lambda x.x$ represent? Intuitively, it's the identity function, but over what domain and codomain? We might think of it as representing the set of all identity functions, but this interpretation quickly leads to Russell's paradox. In fact, defining a precise mathematical model for lambda-calculus terms is far from straightforward, requiring some sophisticated domain theory.

One possible meaning of a term is divergence. There are infinitely many divergent terms; one example is Ω . In some sense, all divergent terms are equivalent, since none of them produce a value. The implication is that it is undecidable to determine whether two terms are equivalent, because otherwise, given the relationship between the λ -calculus and Turing machines, we could solve the halting problem on lambda calculus terms by testing equivalence to Ω .

1.1 Observational equivalence

Another way of approaching the problem is to say that two terms are equivalent if they behave indistinguishably in all possible contexts.

More precisely, two terms will be considered equal if in every context, either

- they both converge and produce the same value, or
- they both diverge.

A *context* is just a term $C[\cdot]$ with a single occurrence of a distinguished special variable, called the *hole*. The notation $C[e]$ denotes the context $C[\cdot]$ with the hole replaced by the term e . Then we then define equality in the following way:

$$e_1 = e_2 \iff \text{for all contexts } C[\cdot], C[e_1] \Downarrow v \text{ iff } C[e_2] \Downarrow v.$$

Without loss of generality, we can simplify the definition to

$$e_1 = e_2 \iff \text{for all contexts } C[\cdot], C[e_1] \Downarrow \text{ iff } C[e_2] \Downarrow,$$

because if they converge to different values, it is possible to devise a context that causes one to converge and the other to diverge. Suppose that $C[e_1] \Downarrow v_1$ and $C[e_2] \Downarrow v_2$, where v_1 and v_2 have different behavior. Then we can find some context $C'[\cdot]$ which applies to v_1 converges, and applied to v_2 diverges. Therefore, the context $C'[C[\cdot]]$ is a context that causes the original e_1, e_2 to converge and diverge respectively, satisfying the simpler definition.

A conservative approximation (but unfortunately still undecidable) is the following. Let e_1 and e_2 be terms, and suppose that e_1 and e_2 converge to the same value when reductions are applied according to some strategy. Then e_1 is equivalent to e_2 . This *normalization* approach (in which terms are reduced to a *normal form* on which no more reductions can be done) is useful for compiler optimization and for checking type equality in some advanced type systems.

2 Rewrite Rules

2.1 Recap— β -reduction

Recall that β -reduction is the following rule:

$$(\lambda x. e) e' \xrightarrow{\beta} e\{e'/x\}.$$

An instance of the left-hand side is called a *redex* and the corresponding instance of the right-hand side is called the *contractum*. For example,

$$\lambda x. \underbrace{(\lambda y. y)x}_{\beta \text{ redex}} \xrightarrow{\beta} \lambda x. x$$

Note that in CBV, $\lambda x. (\lambda y. y) x$ is a value (we cannot apply a β -reduction inside the body of an abstraction) so we cannot apply this reduction.

2.2 Capture-avoiding substitution

CBN and CBV evaluation reduce β redexes $(\lambda x. e) e'$ only when the RHS e' is a closed term. In general we can try to normalize lambda terms by reducing redexes *inside* lambda abstractions, because the reductions should still preserve the equivalence of the terms. However, in this case the term e' may be an open term containing free variables. Substituting e' into e may cause *variable capture*. For example, consider the substitution $(y (\lambda x. x y))\{x/y\}$. If we replace all the unbound y 's with x , we'll get $x (\lambda x. x x)$.

In fact, this is a problem seen in many other mathematical contexts, such as in integral calculus, because like λ , the integral operator is a binder. For example, consider the following naive attempt to evaluate an integral, with a similar variable capture problem:

$$\int_0^x dy \cdot (1 + \int dx \cdot x) = (y + \int dx \cdot yx) \Big|_{y=0}^{y=x} = (x + \int x^2 dx) - 0 = (x + \int x^2 dx)$$

We write $e_1\{e_2/x\}$ to denote the result of substituting e_2 for all free occurrences of x in e_1 according to the following rules:

$$\begin{aligned} x\{e/x\} &= e \\ y\{e/x\} &= y && \text{where } y \neq x \\ (e_1 e_2)\{e/x\} &= e_1\{e/x\} \cdot e_2\{e/x\} \\ (\lambda x. e_0)\{e_1/x\} &= \lambda x. e_0 \\ (\lambda y. e_0)\{e_1/x\} &= \lambda y. e_0\{e_1/x\} && \text{where } y \neq x \text{ and } y \notin FV(e_1) \\ (\lambda y. e_0)\{e_1/x\} &= \lambda z. e_0\{z/y\}\{e_1/x\} && \text{where } y \neq x, z \neq x, z \notin FV(e_0), \text{ and } z \notin FV(e_1). \end{aligned}$$

Note that the rules are applied inductively. That is, the result of a substitution in a compound term is defined in terms of substitutions on its subterms. The very last of the six rules applies when $y \in FV(e_1)$. In this case we can rename the bound variable y to z to avoid capture of the free occurrence of y . One might well ask: But what if y occurs free in the scope of a λz in e_0 ? Wouldn't the z then be captured? The answer is that it will be taken care of in the same way, but inductively on a smaller term.

Despite the importance of substitution, it was not until the mid-1950's that a completely satisfactory definition of substitution was given, by Haskell Curry. Previous mathematicians, from Newton to Hilbert to Church, worked with incomplete or incorrect definitions. It is the last of the rules above that is hard to get right.


In the pure λ -calculus, we can start with a λ -term and perform β -reductions on subterms in any order, using the full substitution rule avoid to avoid variable capture when the substituted term is open.

3 Confluence

In the classical λ -calculus, no reduction strategy is specified, and no restrictions are placed on the order of reductions. Any redex may be chosen to be reduced next. A λ -term in general may have many redexes, so the process is nondeterministic. We can think of a reduction strategy as a mechanism for resolving the nondeterminism, but in the classical λ -calculus, no such strategy is specified. A *value* in this case is just a term containing no redexes. Such a term is said to be in *normal form*.

This makes it very difficult to define equality. One sequence of reductions may terminate, but another may not. It is even conceivable that different terminating reduction sequences result in different values. Luckily, it turns out that the latter cannot happen.

It turns out that the λ -calculus is *confluent* (also known as the *Church–Rosser* property) under α - and β -reductions. Confluence says that if e reduces by some sequence of reductions to e_1 , and if e also reduces by some other sequence of reductions to e_2 , then there exists an e_3 such that both e_1 and e_2 reduce to e_3 .



It follows that up to α -equivalence, normal forms are unique. For if $e \Downarrow v_1$ and $e \Downarrow v_2$, and if v_1 and v_2 are in normal form, then by confluence they must be α -equivalent. Moreover, regardless of the order of previous reductions, it is always possible to get to the unique normal form if it exists.

However, note that it is still possible for a reduction sequence not to terminate, even if the term has a normal form. For example, $(\lambda x. \lambda y. y)\Omega$ has a nonterminating CBV reduction sequence

$$(\lambda xy.y)\Omega \xrightarrow{\beta} (\lambda xy.y)\Omega \xrightarrow{\beta} \dots$$

but a terminating CBN reduction sequence, namely

$$(\lambda x. \lambda y. y)\Omega \xrightarrow{\beta} \lambda y. y.$$

It may be difficult to determine the most efficient way to expedite termination. But even if we get stuck in a loop, the Church–Rosser theorem guarantees that it is always possible to get unstuck, provided the normal form exists.

In *normal order*, the leftmost redex is always reduced first. This strategy is also called *normal order*. This is closely related to CBN evaluation, but also reduces inside lambdas. Like CBN, it finds a normal form if one exists, albeit not necessarily in the most efficient way. Call-by-value (CBV) is correspondingly related to *applicative order*, where arguments are reduced first.

In C, the order of evaluation of arguments is not defined by the language; it is implementation-specific. Because of this, and the fact that C has side effects, C is not confluent. For example, the value of the expression $(x = 1) + x$ is 2 if the left operand of $+$ is evaluated first, $x + 1$ if the right operand is evaluated first. This makes writing correct C programs more challenging!

The absence of confluence in concurrent imperative languages is why concurrent programming is difficult. In the lambda calculus, confluence guarantees that reduction can be done in parallel without fear of changing the result.