

Lecture Notes on Inference Rules and Careful Definition of λ -Calculus

15-814: Types and Programming Languages
Jan Hoffmann and Frank Pfenning

Lecture 4
Thursday, September 3, 2026

1 Introduction

In this lecture, we formalize the definitions that we introduced informally using the concept of *inference rules* and *inductive definitions*. We will use inference rules to define most of the concepts in this course. To prove a lemma about an inductively defined concept, we usually use *rule induction*.

2 Inference Rules and Rule Induction

We now revisit the definition of the λ -calculus to formalize the concepts like β -equality, which were introduced informally. Throughout this course, we use *inference rules* to write down inductive definitions. Inference rules usually define a *judgment*, which can be thought of as a relation between certain objects. An inference rule has the form

$$\frac{J_1 \cdots J_n}{J}$$

where J_i and J are judgments. We call $J_1, \dots, J_n$ the premises of the rule and J the conclusion. The meaning of the rule is that J_1 and J_2 and $\dots J_n$ imply J .

Inference rules can be best understood by example. For instance, we can inductively define a judgment *e exp*, which states that e is an expression. We

define this judgment with 3 inference rules.

$$\frac{x \text{ var}}{x \text{ exp}} E_1 \qquad \frac{x \text{ var} \quad e \text{ exp}}{\lambda x. e \text{ exp}} E_2 \qquad \frac{e_1 \text{ exp} \quad e_2 \text{ exp}}{(e_1 \ e_2) \text{ exp}} E_3$$

Rules are often annotated with names, here E_1 , E_2 , and E_3 , so that we can easily reference them. In the rule E_1 , we assume that we already have a judgment $x \text{ var}$ for variables available. The rule states that each variable is an expression. The other rules reflect the remainder of the inductive definition of expressions.

We can take two views on the meaning of such inductive definitions. The first one is that inference rules define the smallest set of judgments that is closed under the rules. In the example, that is the smallest set that contains $x \text{ exp}$ for each variable, $\lambda x. e \text{ exp}$ if $e \text{ exp}$ is in the set and x is a variable, and $(e_1 \ e_2) \text{ exp}$ if $e_1 \text{ exp}$ and $e_2 \text{ exp}$ are in the set.

Equivalently, we can say that the rules define the set of judgments that are *derivable* using the rules. A judgment is derivable if it has a *derivation tree*.

For example, we can derive $((\lambda x. x) \ y) \text{ exp}$ with the following derivation tree.

$$\frac{\frac{x \text{ var}}{x \text{ exp}} E_1 \quad \frac{y \text{ var}}{y \text{ exp}} E_1}{(\lambda x. x) \text{ exp} \quad y \text{ exp}} E_2 \quad \frac{(\lambda x. x) \text{ exp} \quad y \text{ exp}}{((\lambda x. x) \ y) \text{ exp}} E_3$$

As we have seen, the expressions of the λ -calculus can be described by a grammar that can be seen as an abbreviation of the rules. So it is in general not necessary to write down such inference rules when we define the syntax of a programming language. However, most inductive definitions cannot be defined by a simple grammar. For example, let us define a judgment $\text{size}(e, n)$ that states that an expression e has size n .

$$\frac{x \text{ var}}{\text{size}(x, 1)} S_1 \qquad \frac{x \text{ var} \quad \text{size}(e, n)}{\text{size}(\lambda x. e, n + 1)} S_2 \qquad \frac{\text{size}(e_1, n_1) \quad \text{size}(e_2, n_2)}{\text{size}(e_1 \ e_2, n_1 + n_2 + 1)} S_3$$

We prove a statement about judgments by *rule induction*. This means that for each rule, we assume the statement holds for the premises of the rule (induction hypothesis) and show that it holds for the conclusion. For example, we can prove the following lemma.

Lemma 1 *If $e \text{ exp}$ then there exists an n such that $\text{size}(e, n)$.*

Proof: By induction on the judgment $e \text{ exp}$.

Case E_1 : Then $e = x$ and $x \text{ var}$. But then $\text{size}(x, 1)$ by rule S_1 .

Case E_2 : Then $e = \lambda x.e'$ for some e' , $x \text{ var}$, and $e' \text{ exp}$. We can apply the induction hypothesis to the premise $e' \text{ exp}$ and conclude that there exists n' such that $\text{size}(e', n')$. Then we can apply rule S_2 to derive $\text{size}(\lambda x.e', n' + 1)$.

Case E_3 : Then $e = (e_1 e_2)$ for some e_i , and we have $e_1 \text{ exp}$ and $e_2 \text{ exp}$. So we can apply the induction hypothesis twice and conclude that there exist n_1 and n_2 such that $\text{size}(e_1, n_1)$ and $\text{size}(e_2, n_2)$. But then we can apply the rule S_3 to conclude $\text{size}(e_1 e_2, n_1 + n_2 + 1)$.

□

Again, there are two views we can take to justify rule induction. The first view is that we show that the judgments for which the statement holds are closed under the inference rules. But since the inductively-defined judgments are the smallest such set, we have shown that the judgments are included in the set for which the statement holds. The second view is to think of rule induction as an induction on derivation trees. We make a case distinction on the rule that has been used at the root and assume that the statement holds for the subtrees (induction hypothesis). If you examine the proof of Lemma 1 again, you can even view it as a program that converts a derivation tree of $e \text{ exp}$ into a derivation tree of $\text{size}(e, n)$.

Another statement we could prove is that the size of an expression is unique. Given that we proved Lemma 1 already, we only need to show that, for a given expression e , there is at most one n such that $\text{size}(e, n)$. A good way to formulate the needed lemma is as follows.

Lemma 2 *If $\text{size}(e, n)$ and $\text{size}(e, m)$ then $n = m$.*

Again, we prove the lemma by rule induction. We have two options: We can induct on the judgment $\text{size}(e, n)$ or the judgment $\text{size}(e, m)$, which will not make a difference in the proof because the statement we aim to prove is symmetric.

Let's proceed by induction on $\text{size}(e, n)$. When we come to the case of the Rule S_1 we know $e = x$, $x \text{ var}$, and $n = 1$. But how do we proceed now? We would like to show that $m = 1$. This is intuitively clear because the

only rule that applies to variables is Rule S_1 . We can prove $m = 1$ by an inner rule induction on $\text{size}(x, m)$ (recall that $e = x$). In the case of Rule S_1 , we immediately have $m = 1$. Rule S_2 does not apply because $x \neq (e_1 e_2)$. Similarly, Rule S_3 also does not apply. So it follows that $m = 1$.

The other cases are similar and both require an inner induction on $\text{size}(e, m)$ as well.

In the future, you do not have to write out the inner induction on $\text{size}(e, m)$ but instead simply apply *inversion*, which is stated by the following lemma.

Lemma 3 (Inversion) *Let $\text{size}(e, n)$.*

1. *If $e = \text{var } x$ then $n = 1$.*
2. *If $e = (e_1 e_2)$ then there exists n_1 and n_2 such $\text{size}(e_1, n_1)$, $\text{size}(e_2, n_2)$, and $n = n_1 + n_2 + 1$.*
3. *If $e = \lambda x. e'$ then there exists n such that $\text{size}(e', n')$ and $n = n' + 1$.*

The proof proceeds by induction on $\text{size}(e, n)$ as previously sketched for the case of Rule S_1 .

3 Binding, Scope, and Substitution

We will first discuss abstract binding trees and then return to the λ -calculus to formally define substitution.

3.1 Abstract Binding Trees

In the following, we rigorously define substitution and related concepts like α -equivalence for the λ -calculus. However, these definitions directly generalize to all programming languages we discuss in this course. An economical way to formalize these syntax-related concepts is to abstract from a single programming language and to define them for so-called *abstract binding trees* (ABTs) [Har16]. Then they apply to every language whose syntax is defined using ABTs.

In this course, we are not going to define ABTs but we will use the ABT notation to define the syntax of languages. For example, the syntax of the λ -calculus can be defined as follows.

		Concrete	Abstract
Expressions	$e ::=$	$\lambda x. e$	$\text{lam}(x.e)$
		$e_1 e_2$	$\text{app}(e_1, e_2)$

ABTs are a generalization of abstract syntax trees that add the concepts of variables and binding. In the previous grammar, we do not mention variables explicitly. ABTs automatically come with a notion of variables that range over the different *sorts* of the language. For the λ -calculus, we only have the sort of expressions and variables always stand for expressions. However, other languages have multiple sorts such as expressions and types. Then we have variables that range over expressions and other variables that range over types.

The definition of the syntactic objects (here expressions) is given by a set of *operators*. In the previous example, the operators are `lam` and `app`. Each operator comes with a fixed number of arguments and a binding structure. The binding is indicated by writing $x.e$ like in the lambda abstraction. It means that one variable (x) is bound here and the scope of the binding is e . If multiple variables are bound then we write $x_1, \dots, x_n.e$.

ABTs are useful not only for formally defining general syntactic concepts but also when specifying the syntax and semantics of a language. The advantage is that the scope of bindings and operators is unambiguous. Later in the semester, you will likely prefer specifying languages in ABT notation.

3.2 Free and Bound Variables

The free variables of an expression of the λ -calculus are defined by the judgment $x \in e$, which reads *variable x is free in expression e* .

$$\frac{x \text{ var}}{x \in x} V_1 \quad \frac{y \in e \quad y \neq x}{y \in \text{lam}(x.e)} V_2 \quad \frac{x \in e_1}{x \in \text{app}(e_1, e_2)} V_3 \quad \frac{x \in e_2}{x \in \text{app}(e_1, e_2)} V_4$$

Free variables are placeholders for expressions. Changing the name of a free variable changes the meaning of an expression. This can be seen if the expression is a sub-expression of a larger expression. On the other hand, we can change the name of bound variables without changing the meaning of an expression in any context (as long as we don't introduce name clashes). Since the name of bound variables does not matter, it generally does not make sense to define something like *the bound variables of an expression*.

3.3 Substitution and Alpha-Equivalence

Through substitution we give meaning to variables. Maybe surprisingly, substitution is subtle to formalize. The main difficulty is an issue that

is called *variable capture*. To understand the issue, consider for example the expression $\lambda x.y \ x$, which is a function that applies the variable y to its argument. If we substitute x for y then we get $\lambda x.x \ x$. The result we expected is a function that applies the variable x to its argument but what we got is a function that applies its argument to itself.

One way to solve the problem is to define substitution only if such a variable capture does not happen. We define the judgment $\text{sub}(e_x, x, e, e')$ through the rules below. The meaning is that the expression e' is the result of replacing variable x with expression e_x in expression e .

$$\begin{array}{c}
\frac{}{\text{sub}(e_x, x, x, e_x)} S_1 \qquad \frac{y \neq x}{\text{sub}(e_x, x, y, y)} S_2 \\
\\
\frac{\text{sub}(e_x, x, e_1, e'_1) \quad \text{sub}(e_x, x, e_2, e'_2)}{\text{sub}(e_x, x, \text{app}(e_1, e_2), \text{app}(e'_1, e'_2))} S_3 \\
\\
\frac{}{\text{sub}(e_x, x, \text{lam}(x.e), \text{lam}(x.e))} S_4 \qquad \frac{x \neq y \quad y \notin e_x \quad \text{sub}(e_x, x, e, e')}{\text{sub}(e_x, x, \text{lam}(y.e), \text{lam}(y.e'))} S_5
\end{array}$$

With this approach, substitution is undefined for some expressions: There does not exist an e such that $\text{sub}(x, y, \lambda x.y \ x, e)$. This is however very inconvenient. For example, it is difficult to define β -reduction since we now have to address the cases in which substitution is not defined. Our goal is therefore to define a total version of the judgment sub .

The idea of the total substitution judgment is to use α -equivalence. In future lectures, we will implicitly identify α -equivalent expressions (and types). So if we talk about an expression such as $\lambda x.x$, we mean the α -equivalence class of $\lambda x.x$.

The judgment for α -equivalence is defined by the following rules. Note that the judgment $e_1 =_\alpha e_2$ is indeed an equivalence relation.

$$\begin{array}{c}
\frac{}{e =_\alpha e} A_1 \qquad \frac{e_1 =_\alpha e'_1 \quad e_2 =_\alpha e'_2}{\text{app}(e_1, e_2) =_\alpha \text{app}(e'_1, e'_2)} A_2 \qquad \frac{\text{sub}(y, x, e, e') \quad y \notin e}{\text{lam}(x.e) =_\alpha \text{lam}(y.e')} A_3 \\
\\
\frac{e =_\alpha e'}{e' =_\alpha e} A_4 \qquad \frac{e_1 =_\alpha e_2 \quad e_2 =_\alpha e_3}{e_1 =_\alpha e_3} A_5
\end{array}$$

Equipped with α -equivalence we can now define the result of substitution to rename bound variables to avoid name capturing. The following rule

defines a judgment $\text{sub}_\alpha(e_x, x, e, e')$ that is total in the sense that for every e_x, x and e there exists an e' such that $\text{sub}_\alpha(e_x, x, e, e')$.

$$\frac{\text{sub}(e_x, x, e_0, e') \quad e_0 =_\alpha e}{\text{sub}_\alpha(e_x, x, e, e')} \text{Sub}_\alpha$$

We write $[e_x/x]e$ for (the equivalence class of) an expression e' such that $\text{sub}_\alpha(e_x, x, e, e')$. The substitution $[e_x/x]e$ is well-defined since all such e' are α -equivalent.

4 Beta-Reduction

Finally, we can use inference rules to define β -reduction and β -equality. The rules for β -equality are as follows.

$$\begin{array}{c} \frac{}{\text{app}(\text{lam}(x.e), e') =_\beta [e'/x]e} B_1 \qquad \frac{e_1 =_\beta e'_1}{\text{app}(e_1, e_2) =_\beta \text{app}(e'_1, e_2)} B_{2a} \\[10pt] \frac{e_2 =_\beta e'_2}{\text{app}(e_1, e_2) =_\beta \text{app}(e_1, e'_2)} B_{2b} \qquad \frac{e =_\beta e'}{\text{lam}(x.e) =_\beta \text{lam}(x.e')} B_3 \qquad \frac{}{e =_\beta e} B_4 \\[10pt] \frac{e_1 =_\beta e_2 \quad e_2 =_\beta e_3}{e_1 =_\beta e_3} B_5 \qquad \frac{e =_\beta e' \quad e' =_\beta e}{e =_\beta e} B_6 \end{array}$$

To define β -reduction we would only use rules $B_1, \dots, B_3$. To define the multi-step β -reduction $e \mapsto_\beta^* e'$ we would use only rules $B_1, \dots, B_5$.

5 Limitations of the λ -Calculus

The λ -calculus is arguably the most elegant programming language. If we believe Church's law then it can also express all computable functions. So why should we even consider other programming languages? There are multiple reasons to study other languages and in particular typed languages.

Expressivity One could say that we got more than we bargained for with the λ -calculus. We defined a seemingly simple programming language (which only provides functions) and discovered that it is expressive enough to not only express non-termination but all computable functions and abstractions such as pairs and natural numbers. However, we set out to study

such abstraction in isolation. We also want to study languages that are *normalizing*, that is, languages in which all programs terminate. Both are difficult in the λ -calculus.

Weak Abstractions While we can express abstractions such as pairs and natural numbers in the λ -calculus, these abstractions are not enforced. For example, we can call the function *plus* with arguments that are not Church numerals and the resulting expression can either be evaluated to a normal form or diverges. There will be no warning or error. In this way, the λ -calculus behaves similarly to an assembly language. On the one hand, this makes it difficult to use the λ -calculus to implement complex software. It would simply be hard to debug large λ expressions when we only see some nonsensical result but do not get hints for finding the point at which the computation went wrong. On the other hand, it makes it difficult to reason about programs in the λ -calculus. For example, we cannot prove by induction on the natural numbers that *plus* will always terminate (has a normal form) if its arguments terminate.

Strong Reduction β -reduction and normal (or leftmost-outermost) evaluation differ from evaluation in most programming languages. So far we have considered only evaluation strategies that are based on *strong reduction*. In strong reduction, we allow evaluation of a function without an argument as defined by Rule B_3 .

$$\frac{e =_{\beta} e'}{\text{lam}(x.e) =_{\beta} \text{lam}(x.e')} B_3$$

Evaluation strategies based on *weak reduction* (like call-by-name or call-by-value, which we will discuss in the next lecture) do not allow such evaluations. As a result, all lambda abstractions $\text{lam}(x.e)$ are normal forms.

It would be perfectly possible to define the λ -calculus using a weak reduction strategy. We would still be able to express all computable functions. However, we would have to make some modifications. For instance, the Church numerals would not work in such a setting since, for example, the normal form of the successor of a Church numeral would not be a Church numeral. So we would have to use a different encoding like Barendregt numerals. Moreover, the Church-Rosser theorem would not hold anymore since the order of β -reductions would matter [BLM05] and we would not have the confluence property. Finally, we would only define *a* λ -calculus instead of *the* λ -calculus, which was defined by Church using β -equivalence.

For these reasons, the λ -calculus with a weak reduction strategy loses some of its appeal.

Exercises

Exercise 1 *Prove by rule induction: For all e and n , if $\text{size}(e, n)$ then e exp.*

Exercise 2 *Provide an expression e such that $\text{sub}_\alpha(x, y, \lambda x.y \ x, e)$. Give a derivation tree for the judgment $\text{sub}_\alpha(x, y, \lambda x.y \ x, e)$.*

References

- [BLM05] Tomasz Blanc, Jean-Jacques Lévy, and Luc Maranget. Sharing in the weak lambda-calculus. In Aart Middeldorp, Vincent van Oostrom, Femke van Raamsdonk, and Roel C. de Vrijer, editors, *Processes, Terms and Cycles: Steps on the Road to Infinity, Essays Dedicated to Jan Willem Klop, on the Occasion of His 60th Birthday*, volume 3838 of *Lecture Notes in Computer Science*, pages 70–87. Springer, 2005.
- [Har16] Robert Harper. *Practical Foundations for Programming Languages*. Cambridge University Press, second edition, April 2016.