In this lecture we will

- define induction on a well-founded relation;
- illustrate the definition with some examples, including the inductive definition of free variables $FV(e)$;
- take another look at inference rules.

1 Free Variables Revisited

Recall that some of the substitution rules mentioned the function $FV : \{\lambda\text{-terms}\} \to \mathsf{Var}$ that defines the set of free variables occurring in an expression.

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

Let us examine the definition of the free-variable function FV.

$$\begin{array}{rcl} FV(x) & = & \{x\} \\ FV(e_1\, e_2) & = & FV(e_1) \cup FV(e_2) \\ FV(\lambda x.\, e) & = & FV(e) - \{x\}. \end{array}$$

Why does this definition uniquely determine the function FV? There are two issues here:

- *existence*: whether FV is defined on all λ-terms;
- *uniqueness*: whether the definition is unique.

Of relevance here is the fact that there are three clauses in the definition of FV corresponding to the three clauses in the definition of λ-terms and that a λ-term can be formed in one and only one way by one of these three clauses. Note also that although the symbol FV occurs on the right-hand side in two of these three clauses, they are applied to proper (*proper* = strictly smaller) subterms.

The idea underlying this definition is called *structural induction*. This is an instance of a general induction principle called *induction on a well-founded relation*.

2 Well-Founded Relations

A binary relation $\prec$ is said to be *well-founded* if it has no infinite descending chains. An *infinite descending chain* is an infinite sequence of elements $a_0, a_1, a_2, \ldots$ such that $a_{i+1} \prec a_i$ for all $i \geq 0$. Note that a well-founded relation cannot be reflexive.

Here are some examples of well-founded relations:

- the successor relation $\{(m, m+1) \mid m \in \mathbb{N}\}$ on $\mathbb{N}$;
- the less-than relation $<$ on $\mathbb{N}$;

- the element-of relation $\in$ on sets. The axiom of foundation (or axiom of regularity) of Zermelo–Fraenkel (ZF) set theory implies that $\in$ is well-founded. Among other things, this prevents a set from being a member of itself;
- the proper subset relation $\subset$ on the set of finite subsets of $\mathbb{N}$.

The following are not well-founded relations:

- the predecessor relation $\{(m+1, m) \mid m \in \mathbb{N}\}$ on $\mathbb{N}$ ($0, 1, 2, \ldots$ is an infinite *descending* chain!);
- the greater-than relation $>$ on $\mathbb{N}$;
- the less-than relation $<$ on $\mathbb{Z}$ ($0, -1, -2, \ldots$ is an infinite descending chain);
- the less-than relation $<$ on the real interval $[0, 1]$ ($1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \ldots$ is an infinite descending chain);
- the proper subset relation $\subset$ on subsets of $\mathbb{N}$ ($\mathbb{N}, \mathbb{N} - \{0\}, \mathbb{N} - \{0, 1\}, \ldots$ is an infinite descending chain).

3 Well-Founded Induction

Let $\prec$ be a well-founded binary relation on a set A. Abstractly, a *property* is just a map $P : A \to \mathbb{2}$, or equivalently, a subset $P \subseteq A$ (the set of all $a \in A$ for which $P(a) = \mathit{true}$).

The principle of well-founded induction on the relation $\prec$ says that in order to prove that a property P holds for all elements of A, it suffices to prove that P holds of any $a \in A$ whenever P holds for all $b \prec a$. In other words,

$$\forall a \in A\ (\forall b \in A\ b \prec a \Rightarrow P(b)) \Rightarrow P(a) \quad \Rightarrow \quad \forall a \in A\ P(a). \tag{1}$$

Expressed as a proof rule,

$$\frac{\forall a \in A\ (\forall b \in A\ b \prec a \Rightarrow P(b)) \Rightarrow P(a)}{\forall a \in A\ P(a)}. \tag{2}$$

The basis of the induction is the case when a has no $\prec$-predecessors; in that case, the statement $\forall b \in A\ b \prec a \Rightarrow P(b)$ is vacuously true.

For the well-founded relation $\{(m, m+1) \mid m \in \mathbb{N}\}$, (1) and (2) reduce to the familiar notion of mathematical induction on $\mathbb{N}$: to prove $\forall n\ P(n)$, it suffices to prove that $P(0)$ and that $P(n+1)$ whenever $P(n)$.

For the well-founded relation $<$ on $\mathbb{N}$, (1) and (2) reduce to *strong* induction on $\mathbb{N}$: to prove $\forall n\ P(n)$, it suffices to prove that $P(n)$ whenever $P(0), P(1), \ldots, P(n-1)$. When $n = 0$, the induction hypothesis is vacuously true.

3.1 Equivalence of Well-Foundedness and the Validity of Induction

In fact, one can show that the induction principle (1)–(2) is valid for a binary relation $\prec$ on A if and only if $\prec$ is well-founded.

To show that well-foundedness implies the validity of the induction principle, suppose the induction principle is not valid. Then there exists a property P for which the premise of (2) holds but not the conclusion. Thus P is false for some element $a_0 \in A$. The premise of (2) is equivalent to

$$\forall a \in A\ \neg P(a) \Rightarrow \exists b \in A\ b \prec a \wedge \neg P(b);$$

this implies that there exists an $a_1 \prec a_0$ such that P is false for a_1. Continuing in this fashion, using the axiom of choice one can construct an infinite descending chain $a_0, a_1, a_2, \ldots$ for which P is false, so $\prec$ is not well-founded.

Conversely, suppose that there is an infinite descending chain $a_0, a_1, a_2, \ldots$. Then the property "$a \notin \{a_0, a_1, a_2, \ldots\}$" violates (2), since the premise of (2) holds but not the conclusion.

4 Structural Induction

Now let us define a well-founded relation on the set of all λ-terms. Define $e < e'$ if e is a *proper* subterm of e'. A λ-term e is a *proper* (or *strict*) subterm of e' if it is a subterm of e' and if $e \neq e'$. If we think of λ-terms as finite labeled trees (see Handout A, Lecture 2), then e' is a tree that has e as a subtree. Since these trees are finite, the relation is well-founded. Induction on this relation is called *structural induction*.

We can now show that $FV(e)$ exists and is uniquely defined for any λ-term e. In the grammar for λ-terms, for any e, exactly one case in the definition of FV applies to e, and all references in the definition of FV are to subterms, which are strictly smaller. The function FV exists and is uniquely defined for the base case of the smallest λ-terms $x \in \mathsf{Var}$. So $FV(e)$ exists and is uniquely defined for any λ-term e by induction on the well-founded subexpression relation.

We often have a set of expressions in a language built from a set of *constructors* starting from a set of *generators*. For example, in the case of λ-terms, the generators are the variables $x \in \mathsf{Var}$ and the constructors are the application operator $\cdot$ and the abstraction operators λx. The set of expressions defined by the generators and constructors is the smallest set containing the generators and closed under the constructors.

If a function is defined on expressions in such a way that

- there is one clause in the definition for every generator or constructor pattern,
- the right-hand sides refer to the value of the function only on proper subexpressions,

then the function is well-defined and unique.

5 Inference Rules

We defined small-step and big-step semantics using inference rules. These rules are another kind of inductive definition. To prove properties of them, we would like to use well-founded induction.

To do this, we can change our view and look at reduction as a binary relation. To say that $\langle c, \sigma\rangle \xrightarrow{1} \langle c', \sigma'\rangle$ according to the small-step SOS rules just means that the pair $(\langle c, \sigma\rangle, \langle c', \sigma'\rangle)$ is a member of some reduction relation, which is a subset of $(\mathit{Com} \times \Sigma) \times (\mathit{Com} \times \Sigma)$. In fact, not only is it a relation, it is a partial function.

Here is an example of the kind of the rule we have been looking at so far.

$$\frac{a_1 \xrightarrow{1} a_1'}{a_1 + a_2 \xrightarrow{1} a_1' + a_2} \; (|\, a_1 \,| > 0) \tag{3}$$

Here a_1, a_2, and a_1' are *metavariables*. Everything above the line is part of the *premise*, and everything below the line is the *conclusion*. The expression on the right side is a *side condition*.

A *rule instance* is a substitution for all the metavariables such that the side condition is satisfied. For

example, here is an instance of the above rule:

$$\frac{3 * 4 \xrightarrow{1} 12}{(3 * 4 + 1) \xrightarrow{1} (12 + 1)} (|3 * 4| > 0)$$

where the substitutions are $a_1 = 3 * 4$, $a_1' = 12$, $a_2 = 1$.

With rules like (3), we are usually trying to define some set or relation. For example, this rule might be part of the definition of some reduction relation $\xrightarrow{1}$ that is a subset of $AExp \times AExp$. Such rules are typically of the form

$$\frac{X_1 \ X_2 \ \ldots \ X_n}{X} (\varphi) \tag{4}$$

where $X_1, X_2, \ldots, X_n$ represent elements that are already members of the set or relation being defined, X represents a new member of the relation added by this rule, and φ is a collection of side conditions that must hold in order for the rule to be applied.

The difference between a premise and a side condition is that the side condition is not part of the relation that the rule is trying to define, whereas the premises are. The side condition is merely some restriction that determines when an instance of the rule may be applied.

Now suppose we have written down a set of rules in an attempt to define a set A. How do we know whether A is well-defined? If the rule (4) is in force, then surely we would like to have $X \in A$ whenever $X_1, X_2, \ldots, X_n \in A$ and the side condition φ holds; but this is hardly a definition of A.

6 Set Operators

One approach is to consider inference rules as *monotone set operators*. Suppose we have a rule R of the form (4) specifying a set of rule instances, where X and the X_i are members of some set S. We can view R as a mapping on subsets of S. Given $B \subseteq S$, define

$$R(B) \triangleq \{X \mid \{X_1, X_2, \ldots, X_n\} \subseteq B \text{ and } \frac{X_1 \ X_2 \ \ldots \ X_n}{X} \text{ is an instance of (4)}\}.$$

Then R is a function mapping subsets of S to subsets of S; that is, $R : 2^S \to 2^S$, where 2^S denotes the *powerset* (set of all subsets) of S. An important property of R is that it is *monotone*: if $B \subseteq C$, then $R(B) \subseteq R(C)$.

Now suppose we have a finite set of rules $R_1, \ldots, R_m$. What set $A \subseteq S$ is defined by the rules? At the very least, we would like A to satisfy the following two properties:

- A is *R-consistent*: $A \subseteq R_1(A) \cup \cdots \cup R_n(A)$. We would like this to hold because we would like every element of A to be included in A *only* as a result of applying one of the rules.
- A is *R-closed*: $R_1(A) \cup \cdots \cup R_n(A) \subseteq A$. We would like this to hold because we would like every element that the rules say should be in A actually to be in A.

These two properties together say that $A = R_1(A) \cup \cdots \cup R_n(A)$, or in other words, A should be a *fixpoint* of the set map $\lambda A.\, R_1(A) \cup \cdots \cup R_n(A)$.

There are two natural questions to ask:

- Does this map actually have a fixpoint?
- Is the fixpoint unique? If not, which one should we take?

We will answer these questions next time.