

Introduction

Dependent types are an extension to traditional types (such as the one seen in the simply-typed lambda calculus) to make them more expressive. Specifically, a type system is said to be dependent if it allows types to depend on terms.

This type theory is based on Per Martin-Löf’s original type theory, which was posited as an alternative foundation for mathematics. In it, we can encode any constructive mathematical theorem.

In this assignment, you’ll be implementing a small language which uses dependent types, building up the language to include new features, then ultimately proving some simple theorems using the language. The next few sections define the language formally, and it will be your task to implement the language yourself.

Syntax

$e, \tau ::= x$	Variable reference
\star	The type of types
$(x : \tau_1) \rightarrow \tau_2$	Dependent function types (AKA Π -types)
$\lambda(x : \tau_1). e_2$	Function introduction
$e_1 e_2$	Function elimination
\mathbb{N}	The natural number type
0 succ e	Natural number introduction
elimNat $e_1 e_2 e_3 e_4$	Natural number elimination

A few bits of syntactic sugar:

- We occasionally use the name ‘ $_$ ’ to bind values we don’t care about. This name should not appear in the variable reference position.
- We can elide names in Π -types if they are not bound in the type’s codomain: $e_1 \rightarrow e_2 \equiv (_ : e_1) \rightarrow e_2$.
- We write the function arrow as right-associative: $(x : e_1) \rightarrow (x : e_2) \rightarrow e_3 \equiv (x : e_1) \rightarrow ((x : e_2) \rightarrow e_3)$.
- We write function elimination as left associative: $e_1 e_2 e_3 \equiv (e_1 e_2) e_3$.
- We write a natural number as the repeated application of **succ** to 0. As an example, we say that $4 \equiv \text{succ}(\text{succ}(\text{succ}(\text{succ } 0)))$.

Note: α -equivalence

In the above syntax, we use the non-terminal x to stand for any variable. In general, the particular variable names we choose for a closed term should not matter. For example, $\lambda(x : \mathbb{N}).x \equiv \lambda(y : \mathbb{N}).y$. This property — that consistently renaming a parameter and all of its occurrences does not change the meaning of a term — is called α -equivalence. In the following sections, we often make the implicit assumption that terms are renamed to avoid collision.

Meta-Functions

Environments

Type-checking is performed relative to a **environment**, also called a context, which is often named Γ . An environment is a list of type/variable pairs, written $(x_1 : \tau_1, x_2 : \tau_2, \dots)$. We write $\Gamma(x)$ to mean the type which is most recently associated with the variable x in the environment Γ . Note: this is not necessarily defined for all variables.

$$\Gamma(x) = \begin{cases} \tau & \text{if } \Gamma = (\dots, x : \tau) \\ (\dots)(x) & \text{if } \Gamma = (\dots, x_1 : \tau) \text{ and } x_1 \neq x \end{cases}$$

Free Variables

We say that a variable is **free** in some term when that variable occurs in the term, but is not bound by a binding term (such as a λ or Π term). We write the set of free variables present in a term e as $\mathcal{FV}(e)$.

$$\mathcal{FV}(e) = \begin{cases} \{x\} & \text{if } e = x \\ \mathcal{FV}(\tau_1) \cup (\mathcal{FV}(\tau_2) - \{x\}) & \text{if } e = (x : \tau_1) \rightarrow \tau_2 \\ \mathcal{FV}(\tau_1) \cup (\mathcal{FV}(e_1) - \{x\}) & \text{if } e = \lambda(x : \tau_1).e_1 \\ \mathcal{FV}(e_1) \cup \mathcal{FV}(e_2) & \text{if } e = e_1 e_2 \\ \mathcal{FV}(e_1) & \text{if } e = \text{succ } e_1 \\ \mathcal{FV}(e_1) \cup \mathcal{FV}(e_2) \cup \mathcal{FV}(e_3) \cup \mathcal{FV}(e_4) & \text{if } e = \text{elimNat } e_1 e_2 e_3 e_4 \\ \emptyset & \text{otherwise} \end{cases}$$

Capture-Avoiding Substitution

When evaluating application, we often use a notion of **substitution** - we replace all occurrences of the argument variable with the argument value. This is a suitable mental model for many programs, but a naive notion of substitution can lead to odd bugs. We write $e_1[x \leftarrow e_2]$ to mean “the term e_1 , but with all free occurrences of x replaced with e_2 ”. Note: $e_1[x \leftarrow e_2]$ is not necessarily defined for all terms - some may require rewriting in terms of α -equivalence.

$$e_1[x \leftarrow e_2] = \begin{cases} e_2 & \text{if } e_1 = x \\ (x_1 : \tau_1[x \leftarrow e_2]) \rightarrow \tau_2[x \leftarrow e_2] & \text{if } e_1 = (x_1 : \tau_1) \rightarrow \tau_2 \\ & \text{and } x_1 \notin \{x\} \cup \mathcal{FV}(e_2) \\ \lambda(x_1 : \tau_1[x \leftarrow e_2]).e_3[x \leftarrow e_2] & \text{if } e_1 = \lambda(x_1 : \tau_1).e_3 \\ & \text{and } x_1 \notin \{x\} \cup \mathcal{FV}(e_2) \\ e_3[x \leftarrow e_2] e_4[x \leftarrow e_2] & \text{if } e_1 = e_3 e_4 \\ \text{succ } e_3[x \leftarrow e_2] & \text{if } e_1 = \text{succ } e_3 \\ \text{elimNat } e_3[x \leftarrow e_2] e_4[x \leftarrow e_2] e_5[x \leftarrow e_2] e_6[x \leftarrow e_2] & \text{if } e_1 = \text{elimNat } e_3 e_4 e_5 e_6 \\ e_1 & \text{otherwise} \end{cases}$$

Semantics

There are two relevant judgments, the typing relation and the reduction relation. These two judgments give a semantics to how type-checking and evaluation should work, respectively.

Typing Relation

Our typing judgment takes the form $\boxed{\Gamma \vdash e : \tau}$, read as “environment Γ types e as τ ”.

$$\begin{array}{c}
\frac{\Gamma(x) = \tau}{\Gamma \vdash x : \tau} \text{TYPE-VAR-REF} \qquad \frac{}{\Gamma \vdash \star : \star} \text{TYPE-}\star \qquad \frac{\Gamma \vdash \tau_1 : \star \quad \Gamma, x : \tau_1 \vdash \tau_2 : \star}{\Gamma \vdash (x : \tau_1) \rightarrow \tau_2 : \star} \text{TYPE-II} \\
\\
\frac{\Gamma \vdash \tau_1 : \star \quad \Gamma, x : \tau_1 \vdash \tau_2 : \star \quad \Gamma, x : \tau_1 \vdash e_2 : \tau_2}{\Gamma \vdash \lambda(x : \tau_1).e_2 : (x : \tau_1) \rightarrow \tau_2} \text{TYPE-}\lambda \qquad \frac{\tau_1 \rightsquigarrow^* \tau_2 \quad \Gamma \vdash e : \tau_2}{\Gamma \vdash e : \tau_1} \text{TYPE-EVAL} \\
\\
\frac{\Gamma \vdash e_1 : (x : \tau_1) \rightarrow \tau_2 \quad \Gamma \vdash e_2 : \tau_1}{\Gamma \vdash e_1 e_2 : \tau_2[x \leftarrow e_2]} \text{TYPE-APP} \\
\\
\frac{}{\Gamma \vdash \mathbb{N} : \star} \text{TYPE-N} \qquad \frac{}{\Gamma \vdash 0 : \mathbb{N}} \text{TYPE-0} \qquad \frac{\Gamma \vdash e : \mathbb{N}}{\Gamma \vdash \text{succ } e : \mathbb{N}} \text{TYPE-succ} \\
\\
\frac{\Gamma \vdash e_1 : \mathbb{N} \rightarrow \star \quad \Gamma \vdash e_2 : e_1 0 \quad \Gamma \vdash e_3 : (x : \mathbb{N}) \rightarrow e_1 x \rightarrow e_1 (\text{succ } x) \quad \Gamma \vdash e_4 : \mathbb{N}}{\Gamma \vdash \text{elimNat } e_1 e_2 e_3 e_4 : e_1 e_4} \text{TYPE-elimNat}
\end{array}$$

Example: Polymorphic Identity Function

The following is a derivation which types the polymorphic identity function `id`, defined as follows:

$$\begin{aligned}
&\text{id} : (A : \star) \rightarrow A \rightarrow A \\
&\text{id} = \lambda(A : \star).\lambda(x : A).x
\end{aligned}$$

We can prove that this type is accurate using a derivation:

$$\begin{array}{c}
\frac{\frac{\frac{(A : \star)(A) = \star}{A : \star \vdash A : \star} \text{TYPE-VAR-REF} \quad \frac{\frac{}{\vdash \star : \star} \text{TYPE-}\star \quad \frac{(A : \star, _ : A)(A) = \star}{A : \star, _ : A \vdash A : \star} \text{TYPE-VAR-REF}}{A : \star \vdash A \rightarrow A : \star} \text{TYPE-II}}{\frac{(A : \star)(A) = \star}{A : \star \vdash A : \star} \text{TYPE-VAR-REF}}{\frac{(A : \star, x : A)(A) = \star}{A : \star, x : A \vdash A : \star} \text{TYPE-VAR-REF} \quad \frac{(A : \star, x : A)(x) = A}{A : \star, x : A \vdash x : A} \text{TYPE-VAR-REF}}{A : \star \vdash \lambda(x : A).x : A \rightarrow A} \text{TYPE-}\lambda} \text{TYPE-}\lambda \\
\frac{}{\vdash \lambda(A : \star).\lambda(x : A).x : (A : \star) \rightarrow A \rightarrow A} \text{TYPE-}\lambda
\end{array}$$

Reduction Relation

Our reduction relation takes the form $\boxed{e_1 \rightsquigarrow e_2}$, which can be read as “ e_1 can be reduced to e_2 by one step of reduction”. We often also talk about the transitive closure of this relation, written $\boxed{e_1 \rightsquigarrow^* e_2}$, read as “ e_1 can be repeatedly reduced to e_2 , which can be reduced no further”.

This may seem like a lot of rules, but only the first three do any interesting work (these are called “computational rules”). The rest just thread through reduction to various subterms (these are called “congruence rules”).

$$\frac{}{(\lambda(x : \tau_1).e_1) e_2 \rightsquigarrow e_1[x \leftarrow e_2]} \text{ EVAL-APP}$$

$$\frac{}{\text{elimNat } e_1 e_2 e_3 0 \rightsquigarrow e_2} \text{ EVAL-elimNat-0}$$

$$\frac{}{\text{elimNat } e_1 e_2 e_3 (\text{succ } e_4) \rightsquigarrow e_3 e_4 (\text{elimNat } e_1 e_2 e_3 e_4)} \text{ EVAL-elimNat-succ}$$

$$\frac{\tau_1 \rightsquigarrow \tau_2}{(x : \tau_1) \rightarrow \tau_3 \rightsquigarrow (x : \tau_2) \rightarrow \tau_3} \text{ EVAL-II-DOMAIN}$$

$$\frac{\tau_2 \rightsquigarrow \tau_3}{(x : \tau_1) \rightarrow \tau_2 \rightsquigarrow (x : \tau_1) \rightarrow \tau_3} \text{ EVAL-II-CODOMAIN}$$

$$\frac{e_1 \rightsquigarrow e_2}{\lambda(x : \tau_1).e_1 \rightsquigarrow \lambda(x : \tau_1).e_2} \text{ EVAL-}\lambda\text{-BODY}$$

$$\frac{e_1 \rightsquigarrow e_2}{e_1 e_3 \rightsquigarrow e_2 e_3} \text{ EVAL-RATOR}$$

$$\frac{e_1 \rightsquigarrow e_2}{e_3 e_1 \rightsquigarrow e_3 e_2} \text{ EVAL-RAND}$$

$$\frac{e_1 \rightsquigarrow e_2}{\text{succ } e_1 \rightsquigarrow \text{succ } e_2} \text{ EVAL-succ}$$

$$\frac{e_1 \rightsquigarrow e_5}{\text{elimNat } e_1 e_2 e_3 e_4 \rightsquigarrow \text{elimNat } e_5 e_2 e_3 e_4} \text{ EVAL-elimNat-MOT}$$

$$\frac{e_2 \rightsquigarrow e_5}{\text{elimNat } e_1 e_2 e_3 e_4 \rightsquigarrow \text{elimNat } e_1 e_5 e_3 e_4} \text{ EVAL-elimNat-BASE}$$

$$\frac{e_3 \rightsquigarrow e_5}{\text{elimNat } e_1 e_2 e_3 e_4 \rightsquigarrow \text{elimNat } e_1 e_2 e_5 e_4} \text{ EVAL-elimNat-IND}$$

$$\frac{e_4 \rightsquigarrow e_5}{\text{elimNat } e_1 e_2 e_3 e_4 \rightsquigarrow \text{elimNat } e_1 e_2 e_3 e_5} \text{ EVAL-elimNat-TARGET}$$