The category of simply-typed lambda terms II

" Thus a proof is, not an object, but an act. " — Martin-Löf

These are the notes for the second part of a lecture on the category of simply-typed lambda terms. Recall that for each type $\sigma \in \Phi_\rightarrow$ we have a set Y_σ of variables of that type. As long as this set is infinite, everything works. For a λ-term M the set $FV(M)$ of free variables is a finite (possibly empty) subset of the set

$$Y := \bigcup_{\sigma \in \Phi_\rightarrow} Y_\sigma$$

of typed variables. The existence of free variables works somewhat against our natural inclination to think of a term $M : \sigma \rightarrow \rho$ as a (set-theoretic) <u>function</u> with domain σ. In fact we should instead consider the typing as a <u>constraint</u> (or guarantee) that M will "yield" an output of type ρ on an input of type σ, but where the process involved in this yielding may involve other parameters — the <u>free variables</u> of M.

<u>Example</u> $M = (\lambda x^\sigma . y^\tau)$ has free variable y, and while it yields y on any input of type σ (thus obeying the constraint of its type) this output clearly depends on the value of the parameter y.

We may make these "hidden" dependencies explicit by <u>λ-abstraction</u>, which removes variables from the set of free variables. However we may restore the hidden dependency just as easily (see the next example). This dual process of "making dependencies explicit" and "hiding dependencies" is described by an <u>adjoint pair</u> of functors, and this (together with a colax 2-functor which maps terms M to a subset of their free variables) is our categorical representation of λ-abstraction in the simply typed λ-calculus.

<u>Example</u> With $M = (\lambda x^{\mathfrak{z}}. y^{\mathfrak{J}})$ as above, $N = \lambda y. M$ has type $\mathfrak{J} \to (\mathfrak{z} \to \mathfrak{J})$ and no free variables, while $(N\ y) = ((\lambda y. M)\ y) =_\beta M$ has the free variable restored.

Given a λ-term $M : \rho$ and variable $q : \mathfrak{J}$ there is a commutative diagram in $\mathcal{L}$

$$
\begin{tikzcd}
\mathbb{1} \arrow[rr, "M"] \arrow[dr, "\tilde{M} = \lambda q. M"'] & & \rho \\
& (\mathfrak{J} \to \rho) \arrow[ur, "\mathfrak{U}^q = \lambda u^{\mathfrak{J} \to \rho}. (u\ q)"'] &
\end{tikzcd}
\tag{2.1}
$$

and it is this <u>lifting</u> of M into $\mathfrak{J} \to \rho$ that we wish to express categorically. First we need to understand the set $FV(M)$ which is clearly <u>not</u> an invariant of M under β-reduction.

<u>Example</u> For variables x, y $(a \neq x)$

$$FV(((\lambda a^{\mathfrak{z}}. x^{\mathfrak{J}})\ y^{\mathfrak{z}})) = FV((\lambda a^{\mathfrak{z}}. x^{\mathfrak{J}})) \cup FV(y^{\mathfrak{z}}) = \{x, y\}$$
$$FV(x^{\mathfrak{J}}) = \{x\}.$$

<u>Lemma</u> If $M \twoheadrightarrow_\beta N$ then $FV(N) \subseteq FV(M)$.

<u>Proof</u> In a β-reduction $((\lambda x. P)\ Q) \longrightarrow P[x := Q]$ the free variables on the RHS were either already free in P or Q. $\square$

<u>Defn</u> Given a λ-term M we define

$$FV_\beta(M) := \bigcap_{\substack{\lambda\text{-terms } N \\ \text{with } M =_\beta N}} FV(N)$$

Using that the simply-typed lambda calculus is strongly normalising and confluent, it is equivalent to say that, for $\hat{M}$ the unique normal λ-term in the β-equivalence class of M, we have $FV_\beta(M) = FV(\hat{M})$.

<u>Lemma</u> Given λ-terms $M : \sigma \to \rho$ and $N : \sigma$ we have

$$FV_\beta((M\,N)) \subseteq FV_\beta(M) \cup FV_\beta(N) \qquad (3.1)$$

<u>Proof</u> We may assume M, N normal, in which case there is a chain

$$(M\,N) \twoheadrightarrow_\beta (\widehat{M\,N})$$

whence we are done by the previous lemma. $\square$

<u>Example</u> The inclusion (3.1) may be strict: consider $M = (\lambda a . x)$ $N = y$ as above.

<u>Lemma</u> Given $M : \sigma \to \rho$ and $N : \tau \to \sigma$,

$$FV_\beta(M \circ N) \subseteq FV_\beta(M) \cup FV_\beta(N) \qquad (3.2)$$

<u>Proof</u> Recall $M \circ N := (\lambda x . (M\,(N\,x)))$, $x \notin FV(M) \cup FV(N)$, and apply the same argument. $\square$

<u>Def^n</u> $\mathcal{V}ar$ denotes the lattice of finite subsets of $\mathcal{Y} = \bigcup_{\sigma \in \Phi_\to} \mathcal{Y}_\sigma$.

At this point we may view $FV_\beta(-)$ as a function on $\mathcal{L}$, sending morphisms to elements of the lattice $\mathcal{V}ar$ (by convention $FV_\beta(*) = \phi$), and having a kind of weak functoriality expressed by (3.2).

<u>Remark</u> If we view $\mathcal{L}$ as a 2-category with only identity 2-morphisms and $\mathcal{V}ar$ as a 2-category with one object, and composition of 1-morphisms (= finite subsets of $\mathcal{Y}$) as union, then FV_β is a colax functor $\mathcal{L} \longrightarrow \mathcal{V}ar$.

<u>Def^n</u> For a finite subset $P \subseteq \mathcal{Y}$ define $\mathcal{L}_P \subseteq \mathcal{L}$ to be the subcategory with the same objects as $\mathcal{L}$, and morphisms

$$\mathcal{L}_P(\sigma, \rho) := \{ f \in \mathcal{L}(\sigma, \rho) \mid FV_\beta(f) \subseteq P \}.$$

<u>Lemma</u> $\mathcal{L}_P$ is a subcategory.

<u>Proof</u> Since $FV_\beta(\mathrm{id}_\sigma) = \phi$ for all σ, $\mathcal{L}_P$ contains identities and (3.2) shows $\mathcal{L}_P$ is closed under composition. $\square$

Let $\rho \in \Phi_\to$ and Q be a finite set of variables. If $Q = \phi$ define $Q^*\rho := \rho$. Otherwise choose an ordering $Q = \{ q_1 : J_1, \dots, q_k : J_k \}$ and set

$$Q^*\rho := J_1 \to J_2 \to \cdots \to J_k \to \rho. \qquad (4.1)$$

and define

$$\mathfrak{U}^Q = \lambda u^{Q^*\rho} . (\cdots ((u\, q_1)\, q_2) \cdots q_k) : Q^*\rho \to \rho \qquad (4.2)$$

As discussed earlier $Q^*\rho$ is independent up to isomorphism of the chosen ordering, in a way which is clearly compatible with $\mathfrak{U}^Q$.

Our main theorem is:

<u>Theorem</u> Given a finite subset $Q \subseteq \mathcal{Y}$ the inclusion $I : \mathcal{L}_{Q^c} \longrightarrow \mathcal{L}$ has a right adjoint, given on objects by $\rho \mapsto Q^*\rho$, with $\mathfrak{U}^Q$ as the counit of adjunction.

<u>Proof</u> Let us first make $Q^*: \mathcal{L} \longrightarrow \mathcal{L}$ into a functor. It clearly suffices to define a functor $F_J = J \longrightarrow (-)$, $J \in \Phi_{\to}$. This is defined on objects in the natural way, with $J \longrightarrow \mathbb{1} := \mathbb{1}$. For β, ρ simple types and $M \in \mathcal{L}(\beta, \rho)$, $M \neq *$ define

$$F_J(M) \in \mathcal{L}(J \to \beta, J \to \rho), \quad F_J(M) := \lambda u^{J\to\beta}.\, \lambda q^J.\, (M(u\,q)) \tag{5.1}$$

$\leftarrow$ q the chosen var of type J in Q

where $u \notin FV(M)$. We set $F_J(*) = *$ in all cases, and for $M \in \mathcal{L}(\mathbb{1}, \rho)$

$$F_J(M) \in \mathcal{L}(\mathbb{1}, J \to \rho), \quad F_J(M) := *.$$

$M_i : \rho_i \to \rho_{i+1}$

$$\begin{aligned} & F_J(M_2 \circ M_1) \\ &= \lambda u.\lambda q.\,(M_2 \circ M_1\,(u\,q)) \\ &= \lambda u.\lambda q.\,(\lambda t.(M_2\,(M_1\,t))\,(u q)) \\ &= \lambda u.\lambda q.\,(M_2\,(M_1\,(u\,q))) \end{aligned}$$

$$\begin{aligned} & F_J(M_2) \circ F_J(M_1) \\ &= \lambda t.\,(F_J(M_2)\,(F_J(M_1)\,t)) \\ &= \lambda t.\,(\lambda q.\,(M_2\,((F_J(M_1)\,t)\,q))) \\ &= \lambda t.\lambda q.\,(M_2\,(\underbrace{(\lambda q.(M_1\,(t\,q)))}\,q)) \\ &= \lambda t.\lambda q.\,(M_2\,(M_1\,(t\,q))) \end{aligned}$$

Next we prove that there is a natural bijection, for any $\beta, \rho \in \Phi_{\to}$

$$\Phi : \mathcal{L}(\beta, Q^*\rho)_{Q^c} \longrightarrow \mathcal{L}(\beta, \rho)$$

$$\Phi(N) := \mathfrak{U}^Q \circ N.$$

Naturality in β is clear, and naturality in ρ follows from naturality of $\mathfrak{U}^Q$, which may be ascertained from commutativity of the following for $M \in \mathcal{L}(\beta, \rho)$ and $q : J$

$$\begin{tikzcd} \beta \arrow[r, "M"] & \rho \\ J \to \beta \arrow[u, "\lambda a^{J\to\beta}.(a\,q)"] \arrow[r, "\lambda u^{J\to\beta}.\,\lambda q.\,(M(u q))"'] & J \to \rho \arrow[u, "\lambda b^{J\to\rho}.(b\,q)"'] \end{tikzcd} \tag{5.3}$$

since

$$\begin{aligned} [\lambda b.(b\,q)] \circ [\underbrace{\lambda u.\lambda q\,(M\,(u\,q))}_{N}] &=_\beta \lambda z^{J\to\beta}.\,((N\,z)\,q) \\ &=_\beta \lambda z.\,(\lambda q.\,(M\,(z\,q))\,q) \\ &=_\beta \lambda z.\,(M\,(z\,q)) \\ &=_\beta \lambda z.\,(M\,(\lambda a.(a\,q)\,z)) \end{aligned}$$

We define a function $\Phi' : \mathcal{L}(\beta, \rho) \to \mathcal{L}(\beta, Q^*\rho)_{Q^c}$ as follows:

$$\Phi'(M) := \lambda q_1^{J_1} \ldots \lambda q_k^{J_k}. M \quad : \quad J_1 \to \cdots \to J_k \to \beta \to \rho$$

where we use the isomorphism discussed in Part I $(J_1 \to \cdots \to J_k \to \beta \to \rho) \cong (\beta \to Q^*\rho)$ in $\mathcal{L}$. Clearly $FV_\beta(\Phi'(M)) \subseteq FV(\lambda q_1 \cdots \lambda q_k . M) \subseteq Q^c$.

$\boxed{\Phi \circ \Phi' = id}$ is the statement that for $M \in \mathcal{L}(\beta, \rho)$ the diagram

$$\begin{tikzcd}
\beta \arrow[rr, "M"] \arrow[dr, "\Phi'(M)"'] & & \rho \\
& Q^*\rho \arrow[ur, "\mathcal{U}^Q"'] &
\end{tikzcd} \tag{6.1}$$

commutes, which follows from the calculation (using $\beta\eta$)

$$\begin{aligned}
\mathcal{U}^Q \circ \Phi'(M) &= \lambda t^\beta . (\mathcal{U}^Q (\Phi' M\, t)) \\
&= \lambda t^\beta . (\cdots ((\Phi' M\, t)\, q_1) \cdots q_k) \\
&= \lambda t^\beta . ((\cdots ((\lambda q_1 \cdots \lambda q_k M)\, q_1) \cdots q_k)\, t) \\
&= \cdots \\
&= \lambda t^\beta . (M\, t) \\
&=_\eta M.
\end{aligned} \tag{6.2}$$

$\boxed{\Phi' \circ \Phi = id}$ Suppose $N \in \mathcal{L}(\beta, Q^*\rho)$ with $FV_\beta(N) \cap Q = \phi$. We may assume N itself is normal, if necessary. Then by a calculation like (6.2)

$$\begin{aligned}
\Phi'\Phi(N) &= \lambda q_1 \cdots \lambda q_k . (\mathcal{U}^Q \circ N) = \lambda q_1 \cdots \lambda q_k\, \lambda t . ((\cdots (N\, q_1) \cdots q_k)\, t) \\
&=_\eta \lambda q_1 \cdots \lambda q_k . (\cdots (N\, q_1) \cdots q_k)
\end{aligned}$$

Now we use in an essential way the hypothesis $FV(N) \cap Q = \phi$ to see that by η-equivalence, this is equal in $\mathcal{L}$ to N. $\square$

We have an adjunction

$$\begin{tikzcd} \mathcal{L}_{Q^c} \arrow[r, bend left, "I"] \arrow[r, phantom, "\perp"] & \mathcal{L} \arrow[l, bend left, "Q^*"] \end{tikzcd} \qquad (7.1)$$

with counit $\mathfrak{U}^Q : I \circ Q^* \longrightarrow id$ and unit $\eta : id \longrightarrow Q^* \circ I$ given by

$$\mathcal{L}(\rho, Q^*\rho)_{Q^c} \xrightarrow{\cong} \mathcal{L}(\rho, \rho)$$

That is,

$$\eta_\rho = \Phi'(id_\rho) = \lambda q_1^{\sigma_1} \cdots \lambda q_k^{\sigma_k}. \lambda x^\rho. x^\rho$$

<u>Conclusion</u> λ-abstraction has a unique property: given Q and $M \in \mathcal{L}(\zeta, \rho)$ the λ-abstraction $\lambda q_1 \cdots \lambda q_k . M =: \lambda q\, M$ makes

$$\begin{tikzcd} \zeta \arrow[rr, "M"] \arrow[dr, "\lambda q . M"'] & & \rho \\ & Q^*\rho \arrow[ur, "\mathfrak{U}^Q"'] & \end{tikzcd} \qquad (7.2)$$

commute, and it is the only morphism in $\mathcal{L}_{Q^c}$ with this property. So λ-abstraction is <u>factorisation</u> through the universal morphism $Q^*\rho \longrightarrow \rho$.

<u>Remark</u> If $Q_1 \subseteq Q_2$ then $Q_2^c \subseteq Q_1^c$ so $\mathcal{L}_{Q_2^c} \subseteq \mathcal{L}_{Q_1^c}$.

<u>Example</u> Let R be a commutative ring and $S \subseteq R$ a multiplicately closed set (say $R = \mathbb{C}[x_1,\dots,x_n]$ and $m = (x_1-\lambda_1,\dots,x_n-\lambda_n)$ for some $\lambda_i \in \mathbb{C}$, then $S = R \setminus m$ is mult.-closed).
For an R-module M there is an exact sequence

$$0 \longrightarrow t(M) \longrightarrow M \xrightarrow{\ \psi_M\ } S^{-1}M$$

where ψ_M is canonical. The submodule $t(M)$ of <u>torsion elements</u> defines a functor $t(-) : R\text{-Mod} \longrightarrow \mathcal{A}$ where $\mathcal{A} \subseteq R\text{-Mod}$ is the full subcategory of <u>torsion</u> modules (i.e. N s.t. $S^{-1}N \cong 0$). This t is a right adjoint to the inclusion $\mathcal{A} \hookrightarrow R\text{-Mod}$.

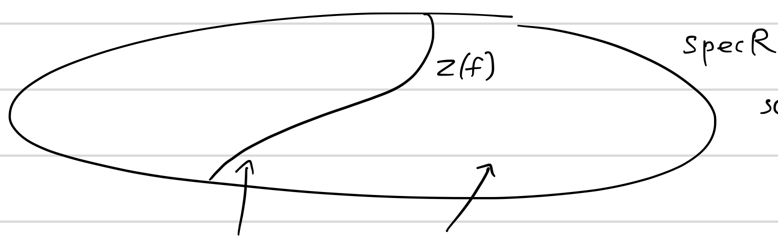

say $S = \{1, f, f^2, \dots\}$
$f \in R$

$Z := \operatorname{Spec}(R/f)$

$t(M)$ are the sections of M supported on Z.

$\operatorname{Spec}(S^{-1}R) =: U$

$S^{-1}M$ is the restriction to the open set U

<u>Def^n</u> A subcategory $\mathcal{A} \hookrightarrow \mathcal{C}$ whose inclusion has a right adjoint is called <u>coreflective</u>.

There are theorems which allow us to recover the <u>space</u> $\operatorname{Spec}(R)$ from the <u>category</u> R-Mod by classifying all reflective / coreflective subcategories (see e.g. Stenström "Rings and modules of fractions").

These ideas lend us the following "geometric" intuition, for what it's worth: a λ-term M is <u>supported</u> on $FV_\beta(M) \subseteq \bar{Y}$ (our Spec R) and if $FV_\beta(M) = \{z_1,\dots,z_n,q\}$ then $M \mapsto \lambda z_1 \cdots \lambda z_n . M$ is analogous to $M \mapsto T_{\{q\}} M$ for modules.