

Topological representation of the λ -calculus

STEVEN AWODEY

*Department of Philosophy, Carnegie Mellon University,
Pittsburgh, PA 15213 USA
Email: awodey@cmu.edu*

Received 5 March 1999; revised 24 September 1999

The λ -calculus can be represented topologically by assigning certain spaces to the types and certain continuous maps to the terms. Using a recent result from category theory, the usual calculus of λ -conversion is shown to be *deductively complete* with respect to such topological semantics. It is also shown to be *functionally complete*, in the sense that there is always a ‘minimal’ topological model in which *every* continuous function is λ -definable. These results subsume earlier ones using cartesian closed categories, as well as those employing so-called Henkin and Kripke λ -models.

1. Introduction

The λ -calculus originated with Church (Church 1940); it is intended as a formal calculus of functional application and specification. In this paper, we are mainly interested in the version known as *simply typed λ -calculus*; as is now well known, the untyped version can be treated as a special case of this (Scott 1980). We present here a topological *representation* of the λ -calculus: types are represented by certain topological spaces and terms by certain continuous functions in such a way that two terms are syntactically equivalent only if they denote the same continuous function, and, moreover, every continuous function between spaces in the representation is denoted by a term.

In particular, then, we are giving *semantics* for the usual deductive system of syntactical ($\beta\eta$) equivalence between terms, with respect to which that system is sound and complete. It is curious that the λ -calculus does not have canonical semantics, but rather exists principally as a syntactical system, about which questions like completeness are always relative to various systems of semantics. In order to compare our results of this kind with previous ones, we will adopt the following terminology:

- A λ -theory \mathbb{T} consists of three sets: (i) a set of basic type symbols; (ii) a set of basic terms, typed over the basic types; (iii) a set of equations between terms in the language $\mathcal{L}[\mathbb{T}]$ over these basic types and terms (see Section 3 below).
- A *model* \mathcal{M} of a theory \mathbb{T} in a category \mathcal{C} (schematically, $\mathcal{M} : \mathbb{T} \rightarrow \mathcal{C}$) interprets the types τ and terms $N(x) : \tau$ of $\mathcal{L}(\mathbb{T})$ as objects and arrows

$$\llbracket N(x) \rrbracket : X \rightarrow \llbracket \tau \rrbracket$$

in \mathcal{C} , in such a way that the equations of \mathbb{T} are satisfied by the interpretations of the

terms equated (Section 4). A model is *standard* if \mathcal{C} is cartesian closed, and function types $\sigma \rightarrow \tau$ are always interpreted as exponentials,

$$\llbracket \sigma \rightarrow \tau \rrbracket = \llbracket \tau \rrbracket^{\llbracket \sigma \rrbracket}.$$

- A system of *semantics* \mathcal{S} for the λ -calculus consists of a class of models $\mathcal{M} : \mathbb{T} \rightarrow \mathcal{C}$, for each theory \mathbb{T} , in possibly different categories \mathcal{C} . By *\mathcal{C} -valued semantics* we mean the collection of models in a fixed category \mathcal{C} . \mathcal{S} is *standard* if all the models in \mathcal{S} are so.
- A system of semantics \mathcal{S} is *complete* if the deductive calculus for syntactic equivalence (Section 3) is sound and complete with respect to \mathcal{S} ; thus, if for every theory \mathbb{T} and any terms $M, N \in \mathcal{L}(\mathbb{T})$,

$$\llbracket M \rrbracket_{\mathcal{M}} = \llbracket N \rrbracket_{\mathcal{M}} \text{ for every model } \mathcal{M} \in \mathcal{S} \quad \text{iff} \quad \mathbb{T} \vdash M = N.$$

\mathcal{S} is *strongly complete* if for every theory \mathbb{T} there is a model \mathcal{M} in \mathcal{S} such that for any terms $M, N \in \mathcal{L}(\mathbb{T})$:

$$\llbracket M \rrbracket_{\mathcal{M}} = \llbracket N \rrbracket_{\mathcal{M}} \quad \text{iff} \quad \mathbb{T} \vdash M = N.$$

Such a model \mathcal{M} is itself said to be *complete*.

- A standard model $\mathcal{M} : \mathbb{T} \rightarrow \mathcal{C}$ is *functionally complete* if every arrow $m : 1 \rightarrow \llbracket \tau \rrbracket$ is of the form $m = \llbracket M \rrbracket$ for some closed term $M : \tau$. In such a model, every arrow $f : \llbracket \sigma \rrbracket \rightarrow \llbracket \tau \rrbracket$ is of the form $f = \llbracket F(x) \rrbracket$ for some term $F(x) : \tau$ with one free variable $x : \sigma$.
- A standard, complete, functionally complete model is called a *representation*. Every arrow $f : \llbracket \sigma \rrbracket \rightarrow \llbracket \tau \rrbracket$ in a representation is of the form $f = \llbracket F(x) \rrbracket$ for a term $F(x) : \tau$ as above, which is, furthermore, unique up to syntactic equivalence modulo \mathbb{T} .

In these terms, the main theorem of this paper (Section 5) states that every λ -theory has a representation in the category of sheaves (equivalently: étale spaces) over a topological space (Section 2). Thus, in particular, *topological semantics*, consisting of all standard models in such categories, are (strongly) complete. It has been an open question for some time whether topological semantics are complete in this sense.

Results of the kind given here go back to L. Henkin (Henkin 1950), who, in effect, showed that non-standard, set-valued semantics are complete (see Mitchell and Scott (1988) for some fine points). An oft-cited result of H. Friedman (Friedman 1975) established the strong completeness of standard, set-valued semantics for the λ -theory consisting of a single basic type (no constants or equations). In the same vein, G. Plotkin has extended the result to certain categories of posets (Mitchell and Moggi 1991). Recently, A. Simpson (Simpson 1996) has shown that the Friedman completeness result is a special case of the following much more general phenomenon: for theories with only basic types, standard \mathcal{C} -valued semantics are complete for any cartesian closed category \mathcal{C} that is not a poset. Moreover, such semantics are strongly complete if \mathcal{C} has a non-repeating endomorphism $a : A \rightarrow A$, that is, such that $a^m = a^n$ implies $m = n$ (Simpson's results rely heavily on the earlier work of R. Statman).

Most theories of interest, however, also involve constants and equations (for example, groups or special data types). Using methods similar to Friedman's, D. Čubrić (Čubrić

1998) has recently extended the (strong) **Set**-valued completeness theorem to theories including constants but no equations. Let us note, however, that this is the end of the line: it is not possible to extend this result further to include arbitrary theories with equations. Indeed, consider the theory of a ‘reflexive domain’: an object D that retracts to its own object of endomorphisms $D \rightarrow D$. This theory can be presented in the form

$$\mathbb{T} = \left\{ \begin{array}{l} D \\ i : (D \rightarrow D) \rightarrow D \\ r : D \rightarrow (D \rightarrow D) \\ \lambda x.ri(x) = \lambda x.x \end{array} \right. \quad (1)$$

with one basic type, two constants, and one equation. (It was introduced by D. Scott to model the untyped λ -calculus (Scott 1980).) Any model of \mathbb{T} in **Set** must have $D = \{*\}$ a singleton set. On the other hand, $\mathbb{T} \not\models \lambda y.ir(y) = \lambda y.y$ (which holds in any such model in **Set**), so **Set**-valued semantics cannot be complete. (A similar argument shows that even semantics in arbitrary *boolean* topoi are not complete.)

We must therefore look beyond **Set** to find complete, standard semantics for the λ -calculus. Of course, the collection of all cartesian closed categories suffices, since every theory has a ‘term model’, which is a representation of it (Section 5). The collection of all categories of the form **Set**^C is then seen to be sufficient, using the well-known Yoneda embedding (*cf.* Section 5 below and Scott (1980)). Finally, W. Mitchell and E. Moggi (Mitchell and Moggi 1991) consider models in categories **Set**^P of functors on *posets* P , which they call ‘Kripke-models’. They establish the strong completeness of *non-standard* semantics in such categories, as a sort of trade-off from standard semantics in arbitrary functor categories **Set**^{C†}.

The topological representation given here thus fits into this line of logical research, implying as it does the strong completeness of standard semantics in categories of sheaves over topological spaces. Since the category of sheaves on a space always has a fully faithful, cartesian closed functor into a category of the form **Set**^P with P a poset, the result of Mitchell and Moggi just cited is here improved upon both by replacing non-standard with standard models, and arbitrary posets with posets of open subsets of a space.

Let us pause briefly to consider the meaning of these various results from a more algebraic point of view. A theory \mathbb{T} generates a free cartesian closed category $\mathcal{C}_{\mathbb{T}}$ with the universal mapping property that models of \mathbb{T} in a cartesian closed category \mathcal{C} correspond uniquely to cartesian closed functors $\mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{C}$, via evaluation at the universal model \mathcal{U} in $\mathcal{C}_{\mathbb{T}}$. Moreover, \mathcal{U} is itself a representation of \mathbb{T} , as can be seen by constructing $\mathcal{C}_{\mathbb{T}}$ ‘syntactically’ so that \mathcal{U} is the term model just mentioned (see Section 5 below). A theory \mathbb{T} then has complete, standard semantics in a cartesian closed category \mathcal{C} if the collection of all cartesian closed functors $\mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{C}$ are jointly faithful, and such \mathcal{C} -valued semantics are strongly (respectively, functionally) complete if there exists a faithful (respectively, full) cartesian closed functor $\mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{C}$. In these terms, then, Friedman’s theorem states that the free cartesian closed category on one object $\mathcal{C}[X]$ has a faithful cartesian closed

† Such non-standard models in categories of functors on posets might better be termed ‘Henkin–Kripke-models’, reserving the term ‘Kripke-models’ for standard models of this form.

functor into **Set**; and the theorem of Čubrić extends this from $\mathcal{C}[X]$ to the free cartesian closed category $\mathcal{C}[G]$ on any graph G . The representation theorem given here asserts the existence of a *full* and faithful, cartesian closed functor,

$$\mathcal{C}_{\mathbb{T}} \longrightarrow \text{sh}(X)$$

into a category of sheaves on a topological space X , for *any* theory \mathbb{T} (and indeed, for any small cartesian closed category \mathcal{C}). This result is analogous to the familiar Freyd embedding theorem for abelian categories (Freyd 1964), and it has an analogous consequence: to show that a diagram in a cartesian closed category commutes, it suffices to assume that the objects are sheaves and the arrows are continuous maps.

The technique of using sheaves on a space as models has its roots in algebraic topology, where, for example, a sheaf with the algebraic structure of a group is regarded as a continuously varying group; such algebraic sheaves are commonly used to obtain models with properties not enjoyed by ‘constant’ algebras. Here we do the same thing, if in somewhat greater generality, by using sheaves of sets to achieve models of conditions not satisfied by constant sets. The algebraic sheaves occurring in topology are in fact a special case of this more general notion; indeed, a sheaf model of the theory of groups in our sense is exactly what the topologist has always meant by a sheaf of groups. In particular, the models considered here are fully determined by the interpretation of the basic language (unlike the so-called ‘Henkin’ models usually considered by logicians). We make this point in order to emphasize that our topological semantics are already well-established in mathematical practice and not cooked-up *ad hoc*.

Before finally getting down to business, let us say a few words about the proof of the main result. It relies essentially on a recent covering theorem for topoi due to C. Butz and I. Moerdijk (Butz 1996; Butz and Moerdijk 1996) (the theorem is related to an earlier one given in Joyal and Moerdijk (1990)). That the main result presented here was not established sooner really attests to the strength of this new covering theorem (rather than to this author’s talents!). Both Butz and Moerdijk have also made essential contributions to the present work.

The contents of this paper are as follows: after summarizing the necessary sheaf theory, in Section 3 we give a brief review of the syntax of the λ -calculus. The reader familiar with either or both of these topics should skip ahead, at most to Section 4 where the topological models at issue are defined. In Section 5 a quick-and-easy proof of the representation theorem is given by using some fairly heavy topos-theoretic machinery. We take the machine apart in Section 6, however, to give an explicit description of the resulting representation in elementary topological terms. The paper concludes by considering extensions of the representation theorem to related logical systems.

2. Sheaves of sets

This section gives an outline of the notions that will be required from the theory of sheaves. Our basic reference is Mac Lane and Moerdijk (1992).

Let X be a fixed topological space. An *étale space* over X is a space E equipped with a *local homeomorphism* $p : E \rightarrow X$; thus p is continuous and every point $e \in E$

has an open neighbourhood $e \in U$ such that $p(U) \subset X$ is open and the restricted map $p|_U : U \rightarrow p(U)$ is a homeomorphism. The fibres $p^{-1}\{x\} \subset E$ for the various points $x \in X$ should be regarded as sets varying continuously over X . A morphism of such étale spaces $(E, p) \rightarrow (E', p')$ is a continuous map $f : E \rightarrow E'$ that is compatible with the structure maps, in the sense that $p'f = p$, as pictured in the commutative triangle:

$$\begin{array}{ccc}
 E & \xrightarrow{f} & E' \\
 & \searrow p & \swarrow p' \\
 & & X
 \end{array} \tag{2}$$

Such a map f is said to be *over* X .

The notion of an étale space over X is equivalent to that of a sheaf on X . Specifically, given an étale space $p : E \rightarrow X$, we often make use of its associated *sheaf of cross-sections*:

$$\Gamma(U, E) = \{s : U \rightarrow E \text{ continuous} \mid ps = i : U \subset X\}.$$

The assignment $U \mapsto \Gamma(U, E)$ is a (contravariant) set-valued functor on the poset of opens $\mathcal{O}(X)$, taking $s \in \Gamma(U, E)$ to the restriction $s|_V \in \Gamma(V, E)$ when $V \subset U$. This functor $\Gamma(-, E)$ satisfies the so-called ‘sheaf’ or ‘patching condition’: given any open set $U \subset X$ with open cover $U = \bigcup_{i \in I} U_i$, and given any cross-sections $s_i \in \Gamma(U_i, E)$ that *match*, in the sense that $s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$ for all $i, j \in I$, there is a unique *amalgamation* $s \in \Gamma(U, E)$, with $s|_{U_i} = s_i$ for each $i \in I$. In general, a *sheaf* on X is just a contravariant, set-valued functor on $\mathcal{O}(X)$ having this same patching property, and a morphism of sheaves is then simply a natural transformation of such functors. It is easy to see that every natural transformation $\Gamma(-, E) \rightarrow \Gamma(-, E')$ between sheaves of sections of étale spaces is induced by a unique continuous map $(E, p) \rightarrow (E', p')$ over X as in 2.

Up to isomorphism of sheaves, *every* sheaf on X is of the form $\Gamma(-, E)$ for some étale space $p : E \rightarrow X$ over X . Indeed, given a sheaf F , one can construct a suitable étale space as follows: let

$$\Lambda F = \sum_{x \in X} F_x, \tag{3}$$

whereby

$$F_x = \varinjlim_{x \in U} F_U. \tag{4}$$

The set F_x is called the *stalk* of F at x . There is an evident projection mapping $p : \Lambda F \rightarrow X$, and we topologize ΛF by using cross-sections of p , as follows. Each stalk $F_x = p^{-1}\{x\} \subset \Lambda F$ consists of so-called *germs* $s_x \in F_x$. By 4 these are equivalence classes of pairs (s, U) with $s \in F_U$, whereby two such (s, U) and (t, V) are deemed equivalent if $s|_W = t|_W$ for some $W \subset U \cap V$. Given any pair (s, U) with $s \in F_U$, taking germs $x \mapsto s_x$ determines a function $\dot{s} : U \rightarrow \Lambda X$ with $p\dot{s}(x) = p(s_x) = x$. We then simply topologize ΛF so that every such function \dot{s} becomes continuous, and is therefore a cross-section of $p : \Lambda F \rightarrow X$. More precisely, the topology on ΛF is generated by basic open sets of the

form $\mathring{s}(U)$ for all $U \in \mathcal{O}(X)$ and $s \in FU$. This clearly also makes $p : \Lambda F \rightarrow X$ continuous, and indeed a local homeomorphism, since given any germ s_x determined by (s, U) , the neighbourhood $\mathring{s}(U)$ is mapped homeomorphically onto the open set U by $p|_{\mathring{s}(U)}$.

We often pass back and forth between the category

$$\text{sh}(X)$$

of all sheaves on X and the category of étale spaces over X , by virtue of the equivalence just outlined. Given a sheaf F , we take its étale space $\Lambda F \rightarrow X$, and given an étale space $E \rightarrow X$, we take its sheaf of sections $\Gamma(-, E)$, making frequent use of the isomorphisms

$$\begin{array}{ll} \Gamma(U, \Lambda F) \cong FU & \text{naturally in } U \in \mathcal{O}(X), \\ \Lambda(\Gamma(-, E)) \cong E & \text{over } X. \end{array}$$

An example of this technique is provided by the specification of the cartesian closed structure on $\text{sh}(X)$. The *product* of two sheaves F and F' on X is simply their product as functors, which is computed ‘pointwise’:

$$(F \times G)(U) = FU \times GU. \tag{5}$$

In terms of étale spaces $p : E \rightarrow X$ and $p' : E' \rightarrow X$, one takes the fibred-product (‘pullback’):

$$E \times_X E' = \sum_{x \in X} E_x \times E'_x$$

as indicated in the diagram

$$\begin{array}{ccc} E \times_X E' & \longrightarrow & E' \\ \downarrow & & \downarrow p' \\ E & \xrightarrow{\quad p \quad} & X \end{array}$$

That these two specifications agree results from the description 4 of the fibre $E_x \times E'_x$ as a stalk of the product sheaf $\Gamma(-, E) \times \Gamma(-, E')$, which is a (filtered) colimit of products 5, hence the product of the stalks. Given any étale space Z over X with maps $f : Z \rightarrow E$ and $f' : Z \rightarrow E'$ over X , the map $(f, f') : Z \rightarrow E \times E'$ over X can then be produced by ‘fibrewise pairing’ $(f, f')_x = (f_x, f'_x) : Z_x \rightarrow E_x \times E'_x$.

Function spaces of sheaves are simply exponentials of functors:

$$G^F(U) = \text{Hom}(yU \times F, G) \tag{6}$$

where $\text{Hom}(-_1, -_2)$ is the set of natural transformations $-_1 \rightarrow -_2$, and yU is the usual contravariant representable functor of the open subset $U \subset X$ (the characteristic function of ‘ $- \subset U$ ’). The function space of étale spaces E and E' can therefore be specified conveniently by applying 6 to the sheaves of sections $\Gamma(-, E)$ and $\Gamma(-, E')$. Of course, the function space $q : E^{E'} \rightarrow X$ can also be described in terms of germs of functions $f : E \rightarrow E'$ over X : the fibre $q^{-1}\{x\}$ over a point $x \in X$ is the set of germs at x of functions $f : E_U \rightarrow E'_U$ over $U \subset X$ with $x \in U$, where $E_U = p^{-1}(U)$ and $p_U : E_U \rightarrow U$ is

the evident restriction of $p : E \rightarrow X$. In these terms, it is easy to describe the evaluation mapping $\epsilon : E^{E'} \times E' \rightarrow E$ over X , and the unique transpose $\tilde{f} : Z \rightarrow E^{E'}$ associated to a map $f : Z \times E' \rightarrow E$ by $\epsilon(\tilde{f}(z), e) = f(z, e)$. We leave these specifications to the reader.

Finally, the constant, singleton-valued sheaf 1 is plainly the sheaf of cross-sections of the identity mapping $1_X : X \rightarrow X$. It is clearly a terminal object, since every étale space $p : E \rightarrow X$ has a unique map $!_E : E \rightarrow X$ over X , namely p itself.

Summing up, we have the following well-known proposition.

Proposition 1. For any topological space X , the category $\text{sh}(X)$ of all sheaves on X with their natural transformations is equivalent to the category of étale spaces and continuous maps over X . It forms a cartesian closed category, with the following canonical structure:

- (i) a terminal object 1 , with a map $!_A : A \rightarrow 1$ for any A ;
- (ii) products $A \times B$, with projections $p_1 : A \times B \rightarrow A$ and $p_2 : A \times B \rightarrow B$ and pairing operation $(f, g) : C \rightarrow A \times B$ for any $f : C \rightarrow A$ and $g : C \rightarrow B$;
- (iii) exponentials B^A with evaluation $\epsilon : B^A \times A \rightarrow B$ and transposition operation $\tilde{f} : C \rightarrow B^A$ for any $f : C \times A \rightarrow B$.

Some readers may wish to know that a *cartesian closed category* is by definition one that has the structure mentioned in (i)–(iii) of the proposition, satisfying the conditions:

$$\begin{aligned} !_A &= f \\ f &= p_1(f, g) \\ g &= p_2(f, g) \\ h &= (p_1 h, p_2 h) \\ f &= \epsilon(\tilde{f} p_1, p_2) \\ g &= (g p_1, p_2) \end{aligned}$$

whenever f , g , and h have the domains and codomains required for these equations to make sense.

3. λ -calculus

This section briefly recalls the standard syntax of the λ -calculus. For a more detailed treatment, the reader can consult Lambek and Scott (1986) and Barendregt (1984).

Types are generated inductively from a set of basic types B, B', \dots by the type-forming operations:

$$\sigma \rightarrow \tau, \sigma \times \tau$$

Terms are built up inductively from *variables* v, v', \dots and a given set of *basic terms* b, b', \dots using the term-forming operations

$$\lambda x : \tau. M, F(M), \langle M, N \rangle, \pi_1(P), \pi_2(P).$$

Term-formation is governed by certain type-restrictions, where the type of a term is a function of the types of the basic terms and variables occurring in it. These type

restrictions and assignments are conveniently stated simultaneously in the rules below, which employ the following conventions:

- A *type assignment* is an expression of the form $M : \tau$, with M a term and τ a type; it can be read ‘ M is a term of type τ ’.
- A *context* (of variables) Γ is a list of type assignments to variables,

$$\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n,$$

with no variable x_i occurring twice.

- A *term in context* is an expression of the form

$$\Gamma \mid M : \tau$$

in which every variable occurring free in the term M also occurs in the context Γ ; it can be read ‘if Γ then $M : \tau$ ’.

The main rules governing term formation are as follows (some rules for manipulating contexts Γ have been omitted):

| | |
|---|--|
| $x : \tau \mid x : \tau$ | $\Gamma \mid M : \tau$ |
| $\Gamma, x : \sigma \mid M : \tau$ | $\Gamma, x : \sigma \mid M : \tau$ |
| $\Gamma \mid (\lambda x : \sigma. M) : (\sigma \rightarrow \tau)$ | $\Gamma \mid M : \sigma \quad \Gamma \mid F : \sigma \rightarrow \tau$ |
| $\Gamma \mid M : \sigma \quad \Gamma \mid N : \tau$ | $\Gamma \mid F(M) : \tau$ |
| $\Gamma \mid \langle M, N \rangle : \sigma \times \tau$ | $\Gamma \mid P : \sigma \times \tau$ |
| | $\Gamma \mid \pi_1(P) : \sigma \quad \Gamma \mid \pi_2(P) : \tau$ |

The type of a closed term is then clearly independent of its context, which we omit when empty.

A *basic language* consists of basic types B, B', \dots and basic terms $b : \tau, b' : \tau', \dots$ (with type assignments). A λ -*theory* \mathbb{T} consists of a basic language and a set of equations $M = N, M' = N', \dots$ between closed terms in the language $\mathcal{L}[\mathbb{T}]$ of all terms over the basic language.

3.1. Syntactic equivalence

Unlike some deductive systems of logic, the λ -calculus is purely equational; we are interested in equations between terms in context. We shall display these in the form

$$\Gamma \mid M = N \tag{7}$$

where Γ is a context of variables including all those occurring free in M or N . Such an equation 7 may be read ‘if Γ , then $M = N$ ’.

The equations and rules that hold for every theory are as follows:

$$\begin{array}{l} \Gamma \mid \langle \pi_1 P, \pi_2 P \rangle = P \\ \Gamma \mid \pi_1 \langle M, N \rangle = M \\ \Gamma \mid \pi_2 \langle M, N \rangle = N \\ \Gamma \mid \lambda x : \tau. (F(x)) = F \quad (x \text{ not free in } F) \\ \Gamma \mid (\lambda x : \tau. N)(M) = N[M/x] \end{array}$$

$$\begin{array}{c} \Gamma \mid M = N \\ \hline \Gamma, x : \tau \mid M = N \end{array} \qquad \begin{array}{c} \Gamma \mid M = M \\ \hline \Gamma \mid M = N \quad \Gamma \mid N = P \\ \hline \Gamma \mid M = P \end{array}$$

$$\begin{array}{c} \Gamma \mid M = N \\ \hline \Gamma \mid N = M \end{array} \qquad \begin{array}{c} \Gamma \mid F = G \quad \Gamma \mid M = N \\ \hline \Gamma \mid F(M) = G(N) \end{array} \qquad \begin{array}{c} \Gamma, x : \tau \mid M = N \\ \hline \Gamma \mid \lambda x : \tau. M = \lambda x : \tau. N \end{array}$$

As usual, the substitution notation $N[M/x]$ is understood to include a convention to prevent binding free variables in M .

Finally, given a theory \mathbb{T} , we define *syntactic equivalence modulo \mathbb{T}* , written

$$\mathbb{T} \vdash M = N,$$

to be the equivalence relation on closed terms $M, N \in \mathcal{L}[\mathbb{T}]$ generated by the equations of \mathbb{T} and the above-stated equations and rules.

4. Topological semantics

We begin by recalling the notion of a model of a λ -theory \mathbb{T} in an arbitrary cartesian closed category \mathcal{C} .

An *interpretation* of the language of \mathbb{T} in \mathcal{C} assigns to each basic type B an object $\llbracket B \rrbracket$ of \mathcal{C} and is then extended inductively to all types by setting

$$\begin{array}{ll} \llbracket \sigma \rightarrow \tau \rrbracket &= \llbracket \tau \rrbracket^{\llbracket \sigma \rrbracket} && \text{(exponential in } \mathcal{C} \text{),} \\ \llbracket \sigma \times \tau \rrbracket &= \llbracket \sigma \rrbracket \times \llbracket \tau \rrbracket && \text{(product in } \mathcal{C} \text{).} \end{array}$$

It also assigns to each basic term $b : \tau$ a morphism $\llbracket b \rrbracket : 1 \rightarrow \llbracket \tau \rrbracket$ of \mathcal{C} . It is extended inductively to any term in context

$$\Gamma \mid M : \tau$$

to yield an arrow

$$\llbracket \Gamma \mid M : \tau \rrbracket : \llbracket \Gamma \rrbracket \longrightarrow \llbracket \tau \rrbracket \tag{8}$$

in \mathcal{C} as follows:

— $\llbracket x : \tau \mid x : \tau \rrbracket = 1_{\llbracket \tau \rrbracket} : \llbracket \tau \rrbracket \rightarrow \llbracket \tau \rrbracket$, for any variable x .

— If $x:\sigma$ does not occur in M ,

$$\llbracket \Gamma, x:\sigma \mid M:\tau \rrbracket = \llbracket \Gamma \mid M:\tau \rrbracket \circ p : \llbracket \Gamma \rrbracket \times \llbracket x:\sigma \rrbracket \rightarrow \llbracket \Gamma \rrbracket \rightarrow \llbracket \tau \rrbracket,$$

where $p : \llbracket \Gamma \rrbracket \times \llbracket x:\sigma \rrbracket \rightarrow \llbracket \Gamma \rrbracket$ is the evident canonical projection.

— $\llbracket \Gamma \mid (\lambda x:\sigma.M):(\sigma \rightarrow \tau) \rrbracket = \llbracket \Gamma, x:\sigma \mid M:\tau \rrbracket$,

where $\widetilde{}$ is the canonical transposition operation.

— $\llbracket \Gamma \mid F(M):\tau \rrbracket = \epsilon \circ \llbracket \Gamma \mid \langle F, M \rangle : (\sigma \rightarrow \tau) \times \sigma \rrbracket$,

where $\epsilon : \llbracket \tau \rrbracket^{\llbracket \sigma \rrbracket} \times \llbracket \sigma \rrbracket \rightarrow \llbracket \tau \rrbracket$ is the canonical evaluation arrow.

— $\llbracket \Gamma \mid \langle M, N \rangle : \sigma \times \tau \rrbracket = (\llbracket \Gamma \mid M:\sigma \rrbracket, \llbracket \Gamma \mid N:\tau \rrbracket)$,

where $(-, -)$ is the canonical pairing operation.

— $\llbracket \Gamma \mid \pi_1(P):\sigma \rrbracket = p_1 \circ \llbracket \Gamma \mid P:\sigma \times \tau \rrbracket$,

where $p_1 : \llbracket \sigma \rrbracket \times \llbracket \tau \rrbracket \rightarrow \llbracket \sigma \rrbracket$ is the first canonical projection.

— $\llbracket \Gamma \mid \pi_2(P):\tau \rrbracket = p_2 \circ \llbracket \Gamma \mid P:\sigma \times \tau \rrbracket$,

where $p_2 : \llbracket \sigma \rrbracket \times \llbracket \tau \rrbracket \rightarrow \llbracket \tau \rrbracket$ is the second canonical projection.

Note that a context Γ is always interpreted by:

$$\begin{aligned} \llbracket \Gamma \rrbracket &= \llbracket x_1:\tau_1, \dots, x_n:\tau_n \rrbracket \\ &= \llbracket \tau_1 \times \dots \times \tau_n \rrbracket \\ &= \llbracket \tau_1 \rrbracket \times \dots \times \llbracket \tau_n \rrbracket. \end{aligned}$$

A closed term $M:\tau$ in the empty context is therefore interpreted as an arrow of the form $1 \rightarrow \llbracket \tau \rrbracket$.

Finally, an interpretation is, of course, a *model* of \mathbb{T} if it satisfies the equations of \mathbb{T} , in the sense that one has identity of arrows $\llbracket M \rrbracket = \llbracket N \rrbracket$ in \mathcal{C} for each equation $M = N$.

Now, a *topological model* of the theory \mathbb{T} is simply a model – in the foregoing sense – in the cartesian closed category of sheaves over a given space X (Section 2). In such a model, each type τ is therefore interpreted as a space E_τ , equipped with a local homeomorphism $p_\tau : E_\tau \rightarrow X$. The operations $\sigma \rightarrow \tau$ and $\sigma \times \tau$ on the types are modelled by the topological product and function-space operations on étale spaces (Section 2). A basic term, and indeed any closed term $c:\tau$, is thus interpreted as a global section $\llbracket c \rrbracket : X \rightarrow E_\tau$ of p_τ , so $p_\tau \llbracket c \rrbracket = 1_X$. In general, a term in context $\Gamma \mid M:\tau$ is interpreted as a continuous function over X , as indicated in the diagram

$$\begin{array}{ccc} E_\Gamma & \xrightarrow{\llbracket \Gamma \mid M:\tau \rrbracket} & E_\tau \\ & \searrow p_\Gamma & \swarrow p_\tau \\ & & X \end{array}$$

The pairing and transposition operations on continuous maps between étale spaces are then used to interpret their syntactic counterparts $\langle -, - \rangle$ and λ in the λ -calculus, in the way just specified. Finally, the syntactic operation $N[M/x]$ of substitution of terms for variables is easily seen to be modelled by composition of continuous functions:

$$\llbracket x:\sigma \mid N:\tau \rrbracket \circ \llbracket \Gamma \mid M:\sigma \rrbracket = \llbracket \Gamma \mid N[M/x]:\tau \rrbracket.$$

5. The main theorem

Theorem 2 (Topological representation). Every λ -theory has a topological representation. Specifically, given a theory \mathbb{T} there is a topological space $X_{\mathbb{T}}$ and a sheaf model $\llbracket \cdot \rrbracket$ over $X_{\mathbb{T}}$ with the following properties:

(i) (*Strong completeness*) For any closed terms N and M ,

$$\llbracket N \rrbracket = \llbracket M \rrbracket \quad \text{iff} \quad \mathbb{T} \vdash N = M.$$

(ii) (*Functional completeness*) For any types σ, τ , every continuous function $f : \llbracket \sigma \rrbracket \rightarrow \llbracket \tau \rrbracket$ over $X_{\mathbb{T}}$ is the interpretation $f = \llbracket x : \sigma \mid F(x) \rrbracket$ of a term $F : \sigma \rightarrow \tau$, which is unique up to syntactic equivalence modulo \mathbb{T} .

The proof proceeds in two main steps:

Step 1. The free cartesian closed category $\mathcal{C}_{\mathbb{T}}$ is characterized uniquely up to isomorphism by the natural (in \mathcal{C}) isomorphism:

$$\text{Hom}(\mathcal{C}_{\mathbb{T}}, \mathcal{C}) \cong \text{Mod}_{\mathbb{T}}(\mathcal{C}), \quad (9)$$

in which $\text{Hom}(\mathcal{C}_{\mathbb{T}}, \mathcal{C})$ denotes the collection of (strict) cartesian closed functors $\mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{C}$, and $\text{Mod}_{\mathbb{T}}(\mathcal{C})$ the collection of \mathbb{T} -models in the cartesian closed category \mathcal{C} . The *universal model* \mathcal{U} is by definition the one associated under 9 to the identity functor $1_{\mathcal{C}_{\mathbb{T}}} : \mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{C}_{\mathbb{T}}$. By naturality of 9, it thus has the property that any \mathbb{T} -model \mathcal{M} in any cartesian closed category \mathcal{C} is the image $\mathcal{M}^*(\mathcal{U}) = \mathcal{M}$ of \mathcal{U} under a unique cartesian closed functor $\mathcal{M}^* : \mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{C}$, as indicated schematically in the diagram

$$\begin{array}{ccc} \mathcal{C}_{\mathbb{T}} & \xrightarrow{\mathcal{M}^*} & \mathcal{C} \\ \mathcal{U} \uparrow & \nearrow \mathcal{M} & \\ \mathbb{T} & & \end{array}$$

Indeed, this universal mapping property may be taken in place of 9 as the characterization of $\mathcal{C}_{\mathbb{T}}$ and \mathcal{U} .

The category $\mathcal{C}_{\mathbb{T}}$ is sometimes called the ‘syntactic category’ and \mathcal{U} the ‘term model’ of \mathbb{T} since they can be constructed from the syntax of \mathbb{T} (Lambek and Scott 1986). Specifically, as objects one takes the types of \mathbb{T} and as arrows $\sigma \rightarrow \tau$ one takes equivalence classes of closed terms $F : \sigma \rightarrow \tau$, identified by syntactic equivalence modulo \mathbb{T} (one must first add a ‘terminal type’ 1 , a new term $* : 1$, and an equation $x : 1 \mid * = x$). The identity maps, composites, and cartesian closed structure are then obvious, as is the ‘identity’ interpretation of \mathbb{T} in $\mathcal{C}_{\mathbb{T}}$. This interpretation is not just a model, however; by virtue of the syntactic definition of the arrows in $\mathcal{C}_{\mathbb{T}}$ it is plainly also complete and functionally complete, which we state as the result of this step in the following lemma.

Lemma 3. Any λ -theory \mathbb{T} has a representation \mathcal{U} in its free cartesian closed category $\mathcal{C}_{\mathbb{T}}$; thus $\llbracket \cdot \rrbracket_{\mathcal{U}}$ satisfies (i) and (ii) of the theorem.

Since we now have all the necessary pieces, we may as well recall the following theorem by the way.

Theorem 4 (Cartesian closed completeness). The collection of models in cartesian closed categories constitutes (strongly) complete semantics.

Proof. In light of the previous lemma, it only remains to show soundness. Given $\mathbb{T} \vdash M = N$ one has $\llbracket M \rrbracket_{\mathcal{U}} = \llbracket N \rrbracket_{\mathcal{U}}$ by the lemma. So for any model $\mathcal{M} = \mathcal{M}^*(\mathcal{U})$ one has $\llbracket M \rrbracket_{\mathcal{M}} = \mathcal{M}^* \llbracket M \rrbracket_{\mathcal{U}} = \mathcal{M}^* \llbracket N \rrbracket_{\mathcal{U}} = \llbracket N \rrbracket_{\mathcal{M}}$. \square

Step 2. For any small cartesian closed category \mathcal{C} , the well-known Yoneda embedding

$$\begin{aligned} y : \mathcal{C} &\longrightarrow \mathbf{Set}^{\mathcal{C}^{\text{op}}} \\ C &\longmapsto \text{hom}_{\mathcal{C}}(-, C) \end{aligned} \quad (10)$$

is full and faithful, and preserves the cartesian closed structure, as the reader can easily check. Since the evaluation functors $\text{eval}_C : \mathbf{Set}^{\mathcal{C}^{\text{op}}} \rightarrow \mathbf{Set}$ for all $C \in \mathcal{C}$ are jointly faithful, we can apply the spatial covering theorem of Butz and Moerdijk (1996) to the topos $\mathbf{Set}^{\mathcal{C}^{\text{op}}}$ to obtain a topological space X and a full and faithful, cartesian closed functor

$$\gamma^* : \mathbf{Set}^{\mathcal{C}^{\text{op}}} \rightarrow \text{sh}(X) \quad (11)$$

into the category of sheaves on X . The functor γ^* is the inverse-image part of a connected, locally-connected geometric morphism $\gamma : \text{sh}(X) \rightarrow \mathbf{Set}^{\mathcal{C}^{\text{op}}}$, the ‘spatial cover’ of the topos $\mathbf{Set}^{\mathcal{C}^{\text{op}}}$. We shall examine the space X more closely in the next section. Composing γ^* with the Yoneda embedding 10 yields a full and faithful, cartesian closed functor $\gamma^*y : \mathcal{C} \rightarrow \text{sh}(X)$, as pictured in the following diagram:

$$\begin{array}{ccc} \mathbf{Set}^{\mathcal{C}^{\text{op}}} & \xrightarrow{\gamma^*} & \text{sh}(X) \\ y \uparrow & \nearrow \gamma^*y & \\ \mathcal{C} & & \end{array}$$

Since the image of a representation under a full and faithful, cartesian closed functor is plainly also a representation, the proof of the theorem is already complete by lemma 3.

6. The topological representation

The foregoing theorem produces for any theory \mathbb{T} a topological space $X_{\mathbb{T}}$ and a complete and functionally complete topological model, comprised of sheaves over $X_{\mathbb{T}}$. Such a model is always given by a full and faithful, cartesian closed functor $\mathcal{C}_{\mathbb{T}} \rightarrow \text{sh}(X_{\mathbb{T}})$ from the free cartesian closed category $\mathcal{C}_{\mathbb{T}}$ (by taking the image of the universal model \mathcal{U} in $\mathcal{C}_{\mathbb{T}}$). In this section, we spell out the spaces and continuous maps involved in this representation of \mathbb{T} . This description is arrived at by applying the construction given in Butz and Moerdijk (1996) to the presheaf topos $\mathbf{Set}^{(\mathcal{C}_{\mathbb{T}})^{\text{op}}}$, taking as a sufficient set of points the evaluation functors

$$\text{eval}_{\sigma} : \mathbf{Set}^{(\mathcal{C}_{\mathbb{T}})^{\text{op}}} \longrightarrow \mathbf{Set}$$

for each object (type) $\sigma \in \mathcal{C}_{\mathbb{T}}$, and then unpacking the result in terms of the λ -theory \mathbb{T} . A similar method was applied to higher-order logic in Awodey and Butz (1997), the appendix of which gives a different perspective on the construction, in terms of so-called ‘Henkin models’.

To simplify the description, we shall assume that the theory \mathbb{T} has countably many basic types and terms, and we add a ‘terminal type’ 1, together with a basic term $* : 1$ and the equation $(x : 1 \mid x = *)$. For each type σ we then choose a distinct variable x_σ and fix the type assignments $x_\sigma : \sigma$ once and for all. If M is a term with at most x_σ free (that is, $\lambda x_\sigma.M$ is closed), we may write $M[x_\sigma]$ to indicate the fact. Finally, put

$$\mathcal{L}[x_\sigma] = \{M \mid \lambda x_\sigma.M \text{ is closed}\}$$

for the set of terms $M[x_\sigma]$.

By an *enumeration* of $\mathcal{L}[x_\sigma]$, we will mean a surjective partial function $f : \mathbb{N} \rightarrow \mathcal{L}[x_\sigma]$ from the natural numbers. Thus f consists of a domain of definition $D_f \subset \mathbb{N}$ and a surjection,

$$\mathbb{N} \supset D_f \xrightarrow{f} \mathcal{L}[x_\sigma].$$

The *points* of the space $X_{\mathbb{T}}$ are equivalence classes of pairs

$$(\sigma, f)$$

where $f : \mathbb{N} \rightarrow \mathcal{L}[x_\sigma]$ is an enumeration of $\mathcal{L}[x_\sigma]$. Two such enumerations (σ, f) and (τ, g) are *equivalent* if $D_f = D_g$ and there exist terms $M[x_\sigma] : \tau$ and $N[x_\tau] : \sigma$ satisfying

$$\begin{aligned} x_\sigma \mid N[M/x_\tau] &= x_\sigma, \\ x_\tau \mid M[N/x_\sigma] &= x_\tau, \\ x_\sigma \mid g_i[M/x_\tau] &= f_i, \\ x_\tau \mid f_i[N/x_\sigma] &= g_i \end{aligned} \tag{12}$$

for all $i \geq 0$. Note, in particular, that two enumerations (σ, f) and (σ, g) are equivalent if $(x_\sigma \mid f_i = g_i)$ for all i , so that each enumeration is equivalent to one (σ, f) with the property that each term $M[x_\sigma]$ has infinitely many ‘labels’ i with $f_i = M[x_\sigma]$ (cf. Butz and Moerdijk (1996)).

To describe the topology on $X_{\mathbb{T}}$, let us say that a term $M[x_\sigma]$ is a *substitution instance* of another $N[x_\tau]$, written

$$M[x_\sigma] < N[x_\tau]$$

if there is a term $S[x_\tau]$ such that

$$x_\sigma \mid M = N[S/x_\tau].$$

The topology of $X_{\mathbb{T}}$ is generated by the following basic open sets: for any $n \geq 1$ and any n -tuples of numbers $(k_i) = (k_1, \dots, k_n)$ and terms $(M_i) = (M_1[x_{\sigma_1}], \dots, M_n[x_{\sigma_n}])$, the set

$$V_{(k_i), (M_i)} = \{(\sigma, f) \mid f(k_i) < M_i[x_{\sigma_i}], i \leq n\} \tag{13}$$

is open. The reader can easily check that these sets $V_{(k_i), (M_i)}$ are indeed well-defined on equivalence classes of enumerations; that is, if (σ, f) and (τ, g) are equivalent, then

$(\sigma, f) \in V_{(k_i), (M_i)}$ iff $(\tau, g) \in V_{(k_i), (M_i)}$. Observe that they are also closed under intersections, since

$$V_{(k_i), (M_i)} \cap V_{(l_j), (N_j)} = V_{(k_i, l_j), (M_i, N_j)}$$

where (k_i, l_j) and (M_i, N_j) denote the evident concatenations of sequences. This completes the description of the space $X_{\mathbb{T}}$.

To describe the representation of \mathbb{T} , let B be a basic type symbol. The space $\llbracket B \rrbracket$ has the underlying point-set

$$\llbracket B \rrbracket = \{(\sigma, f, P) \mid (\sigma, f) \in X_{\mathbb{T}}, P[x_\sigma] : B\},$$

or rather, more precisely, the points of $\llbracket B \rrbracket$ are equivalence classes of such triples, with (σ, f, P) and (τ, g, Q) equivalent if (σ, f) and (τ, g) are equivalent via some terms $M[x_\sigma] : \tau$ and $N[x_\tau] : \sigma$ as before 12, and now also

$$\begin{aligned} x_\tau \mid P[N/x_\sigma] &= Q, \\ x_\sigma \mid Q[M/x_\tau] &= P. \end{aligned} \tag{14}$$

There is an evident projection mapping,

$$p : \llbracket B \rrbracket \rightarrow X_{\mathbb{T}}.$$

The topology on $\llbracket B \rrbracket$ has the basic open sets

$$W_{(k_i), (M_i), m} = \{(\sigma, f, P) \mid (\sigma, f) \in V_{(k_i), (M_i)}, P = f(m)\}, \tag{15}$$

for (k_i) , (M_i) , and $V_{(k_i), (M_i)}$ as in 13, and $m \geq 0$. This topology makes the projection $p : \llbracket B \rrbracket \rightarrow X_{\mathbb{T}}$ a local homeomorphism, and the stalk $p^{-1}\{x\}$ of the resulting sheaf at a point $x = (\sigma, f)$ is then simply

$$\llbracket B \rrbracket_{(\sigma, f)} = \{(\sigma, f, P) \mid P[x_\sigma] : B\}.$$

To determine the interpretation $\llbracket b \rrbracket : 1 \rightarrow \llbracket \tau \rrbracket$ of a basic term $b : \tau$, we must first have the interpretation $p_\tau : \llbracket \tau \rrbracket \rightarrow X_{\mathbb{T}}$ of the type τ . Just as for the basic type B , the fibre $p_\tau^{-1}\{x\}$ over a point $x = (\sigma, f) \in X_{\mathbb{T}}$ is

$$\llbracket \tau \rrbracket_{(\sigma, f)} = \{(\sigma, f, P) \mid P[x_\sigma] : \tau\}.$$

The topology on

$$\llbracket \tau \rrbracket = \sum_{x \in X_{\mathbb{T}}} \llbracket \tau \rrbracket_x$$

then has basic open sets of the same form 15. The interpretation of the basic term $b : \tau$ is to be a global section of the projection mapping $p_\tau : \llbracket \tau \rrbracket \rightarrow X_{\mathbb{T}}$. For this, we of course take the function

$$\begin{aligned} \llbracket b \rrbracket : X_{\mathbb{T}} &\longrightarrow \llbracket \tau \rrbracket, \\ (\sigma, f) &\longmapsto (\sigma, f, b), \end{aligned}$$

which is obviously continuous.

7. Extensions and applications

Topological representations of the kind given here for the simply-typed λ -calculus are also possible for a number of related and richer logical systems, providing these, too, with complete topological semantics. These extensions rest chiefly on the fact that the respective logics (i) have universal models in their syntactic categories and (ii) are preserved by the Yoneda embedding (and possibly sheafification), and by the spatial covering map 11, so that the proof given in Section 5 carries over without substantial change. Rather than going into details, it perhaps suffices to mention the main relevant changes in the various cases.

Untyped λ -calculus.

Following D. Scott the untyped λ -calculus is modelled by a reflexive object in a cartesian closed category, such an object is simply a model of a particular λ -theory (namely 1 above). Indeed, every model of the untyped calculus is logically equivalent to one of this kind (Scott 1980). Since the current topological representation holds for theories with basic terms and equations, applying it to the theory of a reflexive object provides a strongly and functionally complete, topological model for the untyped theory.

Dependent type theory.

The λ -calculus with ‘dependent types’ involves type symbols $\sigma(x)$ containing variables, which themselves may be typed over such dependent types $x : \tau(y)$, and over sum $\sum y : \rho.\tau(y)$ and product $\prod y : \rho.\tau(y)$ types constructed from these. Categorically, such *indexed families* of types are usually modelled using slice categories, and the λ -calculus with dependent types can indeed be modelled in cartesian closed categories having all slice categories also cartesian closed (these are called *locally cartesian closed categories*). The equivalence between simply-typed lambda calculus and cartesian closed categories (Lambek and Scott 1986) extends to λ -calculus with dependent types and locally cartesian closed categories; in particular, theories in this logic have universal models in such categories (*cf.* Seely (1984)). To extend the topological representation to this case, it therefore suffices to observe that locally cartesian closed structure is preserved by both the Yoneda embedding and the spatial covering map (Butz and Moerdijk 1996).

Recently, more elaborate categorical models (involving fibred categories) have been used to model the complex syntax of dependent type theory more closely, and to model systems like P. Martin-Löf’s with additional type theoretic structure (see Pitts (1995) for references). To the extent that such categories are fibrationally equivalent to locally cartesian closed categories, the topological representation still applies. Topological representations of dependent type theory with additional structure are to be treated in (Moerdijk and Palmgren 1999).

Higher-order logic.

Topological representations and completeness theorems for classical and intuitionistic first-order logic have been known (at least to topos-theorists) for some time (Butz and

Johnstone 1998). The methods used here can be applied to give new proofs of these results for first-order logic; they also extend, however, to systems of higher-order logic. For classical logic including propositional connectives and quantification, the topological representation and completeness theorems can be obtained for the full type hierarchy of functions and relations. For the intuitionistic case, the same is true for first order-logic augmented by a hierarchy of function types with quantification at each type, but no relational quantification. Both of these higher-order cases are treated in Awodey and Butz (1997).

References

- Awodey, S. and Butz, C. (1997) Topological completeness for higher-order logic. Tech. Report RS-97-21, BRICS, Aarhus University, Denmark. (Forthcoming in the *Journal of Symbolic Logic*.)
- Barendregt, H. (1984) *The lambda calculus, its syntax and semantics*, North-Holland.
- Butz, C. (1996) *Logical and cohomological aspects of the space of points of a topos*, Ph.D. thesis, Universiteit Utrecht.
- Butz, C. and Johnstone, P. (1998) Classifying toposes for first-order theories. *Annals of Pure and Applied Logic* **91** 33–58.
- Butz, C. and Moerdijk, I. (1996) Topological representation of sheaf cohomology of sites. Tech. Report 973, Universiteit Utrecht.
- Church, A. (1940) A foundation for the simple theory of types. *Journal of Symbolic Logic* **5** 56–68.
- Čubrić, D. (1998) Embedding a free cartesian closed category into the category of sets. *Journal of Pure and Applied Algebra* **126** 121–147.
- Freyd, P. (1964) *Abelian categories: An introduction to the theory of functors*, Harper and Row.
- Friedman, H. (1975) Equality between functionals. In: Parikh, R. (ed.) *Logic Colloquium* **73**, Springer-Verlag.
- Henkin, L. (1950) Completeness in the theory of types. *Journal of Symbolic Logic* **15** 81–91.
- Joyal, A. and Moerdijk, I. (1990) Toposes as homotopy groupoids. *Advances in Mathematics* **80** 22–38.
- Lambek, J. and Scott, P.J. (1986) *Introduction to higher-order categorical logic*, Cambridge University Press.
- Mac Lane, S. and Moerdijk, I. (1992) *Sheaves in geometry and logic: A first introduction to topos theory*, Springer-Verlag.
- Mitchell, J. and Moggi, E. (1991) Kripke-style models for typed lambda calculus. *Annals of Pure and Applied Logic* **51** 99–124.
- Mitchell, J. and Scott, P. (1988) Typed lambda models and cartesian closed categories. *Contemporary Mathematics* **92** 301–316.
- Moerdijk, I. and Palmgren, E. (1999) Wellfounded trees in categories (unpublished manuscript).
- Pitts, A. (1995) Categorical logic. In: Abramsky, A. et al. (eds.) *Handbook of Logic in Computer Science*, vol. VI, Oxford University Press (forthcoming).
- Scott, D. (1980) Relating theories of the lambda calculus. In: Seldin, P. and Hindley, R. (eds.) *To H.B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, Academic Press 403–450.
- Seely, R. (1984) Locally cartesian closed categories and type theory. *Mathematical Proceedings of the Cambridge Philosophical Society* **95** 33–48.
- Simpson, A. (1996) Categorical completeness results for the simply-typed lambda-calculus. Typed Lambda Calculi and Applications, Proceedings of TLCA '95. *Springer-Verlag Lecture Notes in Computer Science* **902** 414–427.