

Type theory and its meaning explanations

Jonathan Sterling

Thanks to Bob Harper, Peter Dybjer, Bengt Nordström, Carlo Angiuli and Darin Morrison for invaluable conversations which helped sharpen my view of the meaning explanations for computational and intensional type theories. All mistakes are mine.

ABSTRACT. At the heart of intuitionistic type theory lies an intuitive semantics called the “meaning explanations”; crucially, when meaning explanations are taken as definitive for type theory, the core notion is no longer “proof” but “verification”. We’ll explore how type theories of this sort arise naturally as enrichments of logical theories with further judgements, and contrast this with modern proof-theoretic type theories which interpret the judgements and proofs of logics, not their propositions and verifications.

Contents

Chapter 1. Logical Theories	5
1. Judgements of a logical theory	5
2. Higher-order judgements	6
3. Propositions and verifications	7
4. Judgements for verifications	8
Chapter 2. Computational Type Theories	11
1. The categorical judgements	11
2. The functional sequent judgement	12
3. The definitions of types	14
Chapter 3. Proof Theories and Logical Frameworks	19
1. Proof-theoretic type theory	19
2. Martin-Löf's equational logical framework	21
3. A critique of MLLF	23
4. A way forward: verifications & uses	25
Bibliography	27

CHAPTER 1

Logical Theories

To start, we will consider the notion of a *logical theory*; in my mind, it begins with a species of judgements that can be proposed, asserted, and (if they are evident) known.

1. Judgements of a logical theory

The basic forms of judgement for a logical theory will be $\boxed{P \text{ prop}}$ and $\boxed{P \text{ true}}$; and what is P ? It is a member of the species of terms, which are made meaningful in the course of making the judgement $P \text{ prop}$ (“ P is a proposition”) evident for a proposition P .

The forms of judgement may be construed as containing *inputs* and *outputs*; an *input* is something which is inspected in the course of knowing a judgement, whereas an *output* is something which is synthesized (or created) in the course of knowing a judgement. The positions of *inputs* and *outputs* in a judgement constitute what is called its *mode*, and we color-code it in this presentation for clarity.¹

To each judgement is assigned a *meaning explanation*, which explicates the knowledge-theoretic content of the judgement. For a judgement \mathcal{J} , a meaning explanation should be in the form:

To know \mathcal{J} is to know...

The meaning of the judgement $P \text{ prop}$ is, then, as follows:

To know $P \text{ prop}$ is to know that P is a proposition, which is to know what would count as a direct verification of P .

So if a symbol P is taken to denote a proposition, we must know *what sort of thing* is to be taken as a direct verification of P , and this is understood as part of the definition of P . A “direct verification” is understood in contrast with an “indirect verification”, which is to be thought of as a means or plan for verifying the proposition; these distinctions will be explained in more detail later on. Now, the judgement $P \text{ true}$ (“ P is true”) is only meaningful in case we know $P \text{ prop}$ (this is called a presupposition). Then the meaning of $P \text{ true}$ is as follows:

¹We will not see any judgements with *outputs* at first, but it will become necessary as soon as we consider judgements about computation, where the reduction of a term is synthesized from the redex. Modes may be used to construe a judgement as having algorithmic content.

To know P true is to have a verification of P .

From the (implicit) presupposition P prop, we already know what counts as a verification, so the meaning explanation is well-defined. Note that having a means or plan for verifying P is equivalent to having a (direct) verification; this follows from the fact that one may put into action a plan for verifying P and achieve such a verification, and likewise, it is possible to propound a plan of verification by appeal to an existing verification.

2. Higher-order judgements

The judgements we have described so far are “categorical” in the sense that they are made without assumption or generality.

2.1. Hypothetical judgement. We will need to define a further form of judgement, which is called “hypothetical”, and this is the judgement under hypothesis $\boxed{\mathcal{J}(\mathcal{J}')}$, pronounced “ \mathcal{J} under the assumption \mathcal{J}' ”. Its meaning explanation is as follows:

To know the judgement $\mathcal{J}(\mathcal{J}')$ is to know the judgement \mathcal{J} assuming you know the judgement \mathcal{J}' .

Hypothetical judgement may be iterated, and $\mathcal{J}(\mathcal{J}_1, \mathcal{J}_2)$ will be used as notation for $\mathcal{J}(\mathcal{J}_2)(\mathcal{J}_1)$.

2.2. General judgement. Another kind of higher order judgement is “general judgement”, which is judgement with respect to a variable, $\boxed{|_x \mathcal{J}}$, pronounced “for an arbitrary x , \mathcal{J} ”. The meaning explanation for this new judgement is as follows:

To know the judgement $|_x \mathcal{J}$ is, to know $[E/x]\mathcal{J}$ (i.e. the substitution of E for x in the expression \mathcal{J}) for any arbitrary expression E ,²

As far as notation is concerned, the bar symbol binds the least tightly of all the other notations we have considered. Likewise, general judgement may be iterated, and the notation $|_{x,y} \mathcal{J}$ will be used as notation for $|_x |_y \mathcal{J}$.

2.3. Hypothetico-general judgement. When hypothetical judgement is used inside general judgement, as in $|_x A(x) \text{ true } (B(x) \text{ true})$, we term the whole thing a “hypothetico-general” judgement. One thing bears clarifying, which is, Why do we write $P \text{ true } (P \text{ true})$ rather than $|_P P \text{ true } (P \text{ true})$?

The former is really not a single judgement, but rather a *scheme* for judgements, where P is intended to be replaced with a concrete expression by the person

²Technically, E is qualified as being of the same valence as x , but because we have not developed a formal theory of expressions in this presentation, I choose to ignore this issue.

asserting the judgements. On the other hand, the latter is itself a single judgement which may be asserted all on its own.

3. Propositions and verifications

Now that we have propounded and explained the minimal system of judgements for a logical theory, let us populate it with propositions. First, we have falsity \perp , and we wish to make $\perp \text{ prop}$ evident; to do this, we simply state what counts as a direct verification of \perp : there is no direct verification of \perp .

The next basic proposition is trivial truth \top , and to make $\top \text{ prop}$ evident, we state that a direct verification of \top is trivial. The definition of \top thus validates the judgement $\top \text{ true}$ (i.e. that we have a verification of \top ; this is immediate).

Next, let us define conjunction; in doing so, we will make evident the hypothetical judgement

$$P \wedge Q \text{ prop } (P \text{ prop}, Q \text{ prop})$$

Equivalently, we can display this as a rule of inference:³

$$\frac{P \text{ prop } \quad Q \text{ prop}}{P \wedge Q \text{ prop}}$$

A direct verification of $P \wedge Q$ consists in a verification of P and a verification of Q ; this validates the assertion of the judgement $P \wedge Q \text{ true}$ ($P \text{ true}, Q \text{ true}$). Because it is a valid inference, we can write it as an inference rule:

$$\frac{P \text{ true } \quad Q \text{ true}}{P \wedge Q \text{ true}}$$

A direct verification of $P \vee Q$ may be got either from a verification of P or one of Q . From this definition we know $P \vee Q \text{ prop}$ ($P \text{ prop}, Q \text{ prop}$), or

$$\frac{P \text{ prop } \quad Q \text{ prop}}{P \vee Q \text{ prop}}$$

The verification conditions of disjunction give rise to two evident judgements $P \vee Q \text{ true}$ ($P \text{ true}$) and $P \vee Q \text{ true}$ ($Q \text{ true}$), which we can write as inference rules:

$$\frac{P \text{ true}}{P \vee Q \text{ true}} \quad \frac{Q \text{ true}}{P \vee Q \text{ true}}$$

Finally, we must define the circumstances under which $P \supset Q$ is a proposition (i.e. when $P \supset Q \text{ prop}$ is evident). And we intend this to be under the circumstances

³Evident hypothetical judgements are often written as rules, i.e.

$$\frac{\text{premise}}{\text{conclusion}}$$

rather than *conclusion* (*premise*). It must be stressed that only *evident/known* judgements may be written in this way.

that P is a proposition, and also that Q is a proposition assuming that P is true. In other words, $P \supset Q \text{ prop}$ ($P \text{ prop}, Q \text{ prop}$ ($P \text{ true}$)), or

$$\frac{P \text{ prop} \quad Q \text{ prop} (P \text{ true})}{P \supset Q \text{ prop}}$$

Now, to validate this judgement will be a bit more complicated than the previous ones. But by unfolding the meaning explanations for hypothetical judgement, proposition-hood and truth of a proposition, we arrive at the following explanation:

To know $P \supset Q \text{ prop}$ ($P \text{ prop}, Q \text{ prop}$ ($P \text{ true}$)) is to know what counts as a direct verification of $P \supset Q$ when one knows what counts as a direct verification of P , and, when one has such a verification, what counts as a direct verification of Q .⁴

If the judgement $P \supset Q \text{ prop}$ ($P \text{ prop}, Q \text{ prop}$ ($P \text{ true}$)) is going to be made evident, then we must come up with what should count as a direct verification of $P \supset Q$ under the assumptions described above.

And so to have a direct verification of $P \supset Q$ is to have a verification of Q assuming that one has one of P ; this is the meaning of implication, and it validates the judgement $P \supset Q \text{ true}$ ($Q \text{ true}$ ($P \text{ true}$)), and may be written as an inference rule as follows:

$$\frac{Q \text{ true} (P \text{ true})}{P \supset Q \text{ true}}$$

4. Judgements for verifications

So far, we have given judgements which define what it means to be a proposition, namely $P \text{ prop}$, and thence for each proposition, we have by definition a notion of what should count as a verification of that proposition. And we have a judgement $P \text{ true}$, which in its assertion means that one has (a way to obtain) such a verification of P , but we have not considered any judgements which actually refer to the verifications themselves symbolically.

It is a hallmark of Martin-Löf’s program to resolve the contradiction between syntax and semantics not by choosing symbols over meanings or meanings over symbols, but by endowing symbols with meaning in the course of knowing the evident judgements. **As such, P is a symbol, but when we assert $P \text{ prop}$ we are saying that we know what proposition P denotes.**

A similar thing can be done with verifications themselves, by representing them with symbols in the same way we have done for the propositions. And then, we can consider a judgement such as “ M is a verification of P ”, and in knowing that

⁴Note that unless $P \text{ true}$, it need not be evident that $Q \text{ prop}$; in other words, Q only has to be a proposition if P is true. It would also be acceptable to give a stronger definition to implication, but this is the one accepted by Martin-Löf.

judgement, we know what verification M is meant to denote. In practice, this judgement has been written in several ways:

Notation	Pronunciation
$M \in P$	M is an element of P
$M \Vdash P$	M realizes P
$P \sqsubseteq_{\text{ext}} M \sqsubseteq$	P is witnessed by M

But they all mean the same thing, so we will choose the notation $\boxed{M \in P}$ and pronounce it “ M verifies P ”. Tentatively, the following defective meaning explanation could be given:

* To know $M \in P$ is to know that M is a verification of P .

But now that we have started to assign expressions to verifications, we must be more careful about differentiating *direct verifications* (which we will call “canonical”) from *indirect verifications* (which we will call “non-canonical”). So the domain of expressions must itself be accorded with a notion of reduction to canonical form, and this corresponds with putting into action a plan of verification in order to get a direct (canonical) verification; reduction to canonical form will be represented by a judgement $\boxed{M \Downarrow M'}$, pronounced “ M evaluates to M' ”.

To know $M \Downarrow M'$ is to know that M is an expression which reduces to a canonical form M' .

An example of an evident reduction judgement in elementary mathematics would be $3 + 4 \Downarrow 7$; note that $3 + 4 \Downarrow 1 + 6$ is, on the other hand, not evident, since this judgement describes reduction to *canonical* form, whereas $1 + 6$ is not a canonical number.

Now, we can correct the previous meaning explanation as follows:

To know $M \in P$ is to know an M' such that $M \Downarrow M'$ and M' is a canonical (direct) verification of P .

If it is not yet clear why it would have been a mistake to fail to use the notion of reduction to canonical form in the above meaning explanation, consider that each time a proposition is defined, it should be possible to do so without knowing what other propositions exist in the theory. But if we consider non-canonical forms (as would be necessary if we omitted the $M \Downarrow M'$ premise), then we would have to fix in advance all the possible non-canonical forms in the computation system in the course of defining each proposition. As such, the open-ended nature of the logic would be destroyed; in a later chapter, the seriousness of this problem will be made even more clear.

The meaning explanation for P prop must be accordingly modified to take into account the computational behavior of expressions:

To know $P \text{ prop}$ is to know a P' such that $P \Downarrow P'$ and P' is a canonical proposition, which is to say, that one knows what counts as a canonical verification for P' .

In practice, when it is clear that P is canonical, then we will simply say, “To know $P \text{ prop}$ is to know what counts as a canonical verification of P ”. As an example, then, we will update the evidence of the following assertion:

$$P \supset Q \text{ prop } (P \text{ prop}, Q \text{ prop } (P \text{ true}))$$

The meaning of this, expanded into spoken language, is as follows:

To know $P \supset Q \text{ prop } (P \text{ prop}, Q \text{ prop } (P \text{ true}))$ is to know what counts as a canonical (direct) verification of $P \supset Q$ under the circumstances that $P \Downarrow P'$, such that one knows what counts as a canonical verification P' , and, if one has such a verification, $Q \Downarrow Q'$ such that one knows what counts as a canonical verification of Q' .

And the above judgement is evident, since we will say that a canonical verification of $P \supset Q$ is an expression $\lambda x.E$ such that we know the hypothetico-general judgement $|_x E \in Q (x \in P)$. This validates the assertion $\lambda x.E \in P \supset Q (|_x E \in Q (x \in P))$, or, written as an inference rule:

$$\frac{|_x E \in Q (x \in P)}{\lambda x.E \in P \supset Q}$$

By the addition of this judgement, we have graduated from a logical theory to a type theory, in the sense of *Constructive Mathematics and Computer Programming* (Martin-Löf, 1979). In fact, we may dispense with the original $P \text{ true}$ form of judgement by *defining* it in terms of the new $M \in P$ judgement as follows:

$$\frac{M \in P}{P \text{ true}}$$

Computational Type Theories

As alluded to at the end of the previous chapter, we may add a judgement $M \in A$ which deals directly with the objects M which verify the propositions (or types) P . We will develop a full theory of dependent types in the sense of Martin-Löf.¹

1. The categorical judgements

Because we will need to consider the introduction of types which do not have a trivial (intensional) equality relation, we must first amend the meaning explanations for some of our judgements, and add a few new forms of judgement. First, we will refer to *types* rather than *propositions* in order to emphasize the generality of the theory; in some presentations, the word *set* is used instead.

The meaning of hypothetical and general judgement are the same as in the previous chapter, and so we will not reproduce them here. The first form of judgement is $\boxed{A \text{ type}}$, and its meaning explanation is as follows:

To know $A \text{ type}$ is to know an A' such that $A \Downarrow A'$, and you know what counts as a canonical verification of A' and when two such verifications are equal.

The next form of judgement is $\boxed{M \in A}$, which remains the same as before:

To know $M \in A$ (presupposing $A \text{ type}$) is to know an M' such that $M \Downarrow M'$ and M' is a canonical verification of A .

We'll need to add new judgements for equality (equality of verifications, and equality of types respectively). First, equality of verifications is written $\boxed{M = N \in A}$, and means the following:

To know $M = N \in A$ (presupposing $M \in A$ and $N \in A$, and thence $A \text{ type}$) is to know that the values of M and N are equal as canonical verifications of A .

¹Please note that the judgements given here, and their meaning explanations, are *not* the same as those used in Constable et al's "Computational Type Theory" and Nuprl. In this chapter, we use the term "computational type theory" in a general sense to characterize a family of type theories which have their origin in Martin-Löf's 1979 paper *Constructive Mathematics and Computer Programming*.

The fact that M and N reduce to canonical values which verify A is known from the presuppositions of the judgement; and what it means for them to be equal as such is known from the evidence of the presupposition A type which is obtained from the other presuppositions.

For equality of types $\boxed{A = B \text{ type}}$, there are a number of possible meaning explanations, but we'll use the one that Martin-Löf used starting in 1979:

To know $A = B$ type (presupposing A type and B type) is to know $|_M M \in B$ ($M \in A$) and $|_M M \in A$ ($M \in B$), and moreover $|_{M,N} M = N \in B$ ($M = N \in A$) and $|_{M,N} M = N \in A$ ($M = N \in B$).

In other words, two types are equal when they have the same canonical verifications, and moreover, the same equality relation over their canonical verifications. Note that there are other possible explanations for type equality, including more intensional ones that appeal to the syntactic structure of type expressions, and these turn out to be more useful in proof assistants for practical reasons. However, the extensional equality that we have expounded is the easiest and most obvious one to formulate, so we will use it here.

2. The functional sequent judgement

Now, because we are allowing the definition of types with arbitrary equivalence relations, we cannot use plain hypothetico-general judgement in the course of defining our types. For instance, if we were going to try and define the function type $A \supset B$ in the same way as we did in the previous chapter, we would permit “functions” which are not in fact functional, i.e. they do not take equal inputs to equal outputs. As such, we will need to bake functionality (also called extensionality) into the definition of functions, and since we will need this in many other places, we elect to simplify our definitions by baking it into a single judgement which is meant to be used instead of plain hypothetico-general judgement.

The judgement which expresses simultaneously generality, hypothesis and functionality has been written in multiple ways. Martin-Löf has always written it as $\mathcal{J}(\Gamma)$, but this is a confusing notation because it appears as though it is merely a hypothetical judgement (but it is much more, as will be seen). Very frequently, it is written with a turnstile, $\Gamma \vdash \mathcal{J}$, and in the early literature surrounding Constable's Computational Type Theory and Nuprl, it was written $\boxed{\Gamma \gg \mathcal{J}}$; we choose this last option to avoid confusion with a similar judgement form which appears in proof-theoretic, intensional type theories; we'll call the judgement form a “(functional) sequent”.

We will define several forms of judgment simultaneously with the sequent judgments: $\boxed{\Gamma \text{ ctx}}$ (“ Γ is a context”), $\boxed{x \# \Gamma}$ (“ x is fresh in Γ ”).

To know $\Gamma \text{ ctx}$ is to know that $\Gamma \equiv \cdot$, or it is to know a variable x and expressions Δ, A such that $\Gamma \equiv \Delta, x : A$ and $\Delta \gg A \text{ type}$, and $x \# \Delta$.

To know $x \# \Gamma$ (presupposing $\Gamma \text{ ctx}$) is to know that $\Gamma \equiv \cdot$, or, if $\Gamma \equiv \Delta, y : A$ such that x is not y and $x \# \Delta$.

In other words, the well-formed contexts are inductively generated by the following grammar:

$$\frac{}{\cdot \text{ ctx}} \quad \frac{\Gamma \text{ ctx} \quad \Gamma \gg A \text{ type} \quad x \# \Gamma}{\Gamma, x : A \text{ ctx}} \\ \frac{}{x \# \cdot} \quad \frac{x \# \Gamma}{x \# \Gamma, y : A}$$

Next, we define the notion of an *environment* or *substitution* for a context, $\boxed{\rho \in^* \Gamma}$. It is possible to define this more generally as functional maps from one context to another, but for our purposes, this is the only case we need.

To know $\rho \in^* \Gamma$, presupposing $\Gamma \text{ ctx}$, is to know, if $\Gamma \equiv \cdot$, then $\rho \equiv \cdot$; otherwise, if $\Gamma \equiv \Delta, x : A$ then $\rho \equiv \rho', M$ such that $\rho' \in^* \Delta$ and $[\rho'/\Delta]M \in [\rho'/\Delta]A$. The substitution $[\rho/\Gamma]$ is effected by replacing each variable in Γ with the corresponding term in ρ .

We will also specify when two environments are equal, $\boxed{\rho_0 = \rho_1 \in^* \Gamma}$:

To know $\rho_0 = \rho_1 \in^* \Gamma$, presupposing $\Gamma \text{ ctx}$, $\rho_0 \in^* \Gamma$ and $\rho_1 \in^* \Gamma$ is to know, if $\Gamma \equiv \cdot$, then $\rho_0 \equiv \cdot$ and $\rho_1 \equiv \cdot$; otherwise, if $\Gamma \equiv \Delta, x : A$ and thence $\rho_0 \equiv \rho'_0, M_0$ and $\rho_1 \equiv \rho'_1, M_1$, then $\rho'_0 = \rho'_1 \in^* \Delta$ and $[\rho_0/\Delta]M_0 = [\rho_1/\Delta]M_1 \in [\rho_0/\Delta]A$.

We will say that $\Gamma \gg \mathcal{J}$ is only a judgement under the presuppositions that $\Gamma \text{ ctx}$ and that \mathcal{J} is a categorical judgement of the form $A \text{ type}$, $A = B \text{ type}$, $M \in A$, or $M = N \in A$. Now we may begin giving the meaning explanations for $\Gamma \gg \mathcal{J}$, starting with typehood with respect to a context:

To know $\Gamma \gg A \text{ type}$ (presupposing $\Gamma \text{ ctx}$) to know $|_\rho [\rho/\Gamma]A \text{ type}$ ($\rho \in^* \Gamma$) and moreover, to know $|_{\rho_0, \rho_1} [\rho_0/\Gamma]A = [\rho_1/\Gamma]A \text{ type}$ ($\rho_0 = \rho_1 \in^* \Gamma$).

We can explain type equality sequents in a similar way:

To know $\Gamma \gg A = B \text{ type}$ (presupposing $\Gamma \text{ ctx}$, $\Gamma \gg A \text{ type}$, $\Gamma \gg B \text{ type}$) is to know

$$|_\rho [\rho/\Gamma]A = [\rho/\Gamma]B \text{ type} \quad (\rho \in^* \Gamma)$$

and moreover, to know

$$|_{\rho_0, \rho_1} [\rho_0/\Gamma]A = [\rho_1/\Gamma]B \text{ type } (\rho_0 = \rho_1 \in^* \Gamma)$$

Next, the meaning of membership sequents is explained:

To know $\Gamma \gg M \in A$ (presupposing $\Gamma \text{ ctx}$, $\Gamma \gg A \text{ type}$) is to know $|_{\rho} [\rho/\Gamma]M \in [\rho/\Gamma]A$ ($\rho \in^* \Gamma$) and moreover, to know

$$|_{\rho_0, \rho_1} [\rho_0/\Gamma]M = [\rho_1/\Gamma]M \in [\rho_0/\Gamma]A \text{ } (\rho_0 = \rho_1 \in^* \Gamma)$$

Finally, member equality sequents have an analogous explanation:

To know $\Gamma \gg M = M' \in A$ (presupposing $\Gamma \text{ ctx}$, $\Gamma \gg A \text{ type}$) is to know $|_{\rho} [\rho/\Gamma]M = [\rho/\Gamma]N \in [\rho/\Gamma]A$ ($\rho \in^* \Gamma$) and moreover, to know

$$|_{\rho_0, \rho_1} [\rho_0/\Gamma]M = [\rho_1/\Gamma]N \in [\rho_0/\Gamma]A \text{ } (\rho_0 = \rho_1 \in^* \Gamma)$$

The simultaneous definition of multiple judgements may seem at first concerning, but it can be shown to be non-circular by induction on the length of the context Γ .

3. The definitions of types

We will now define the types of a simple computational type theory without universes. In the course of doing so, opportunities will arise for further clarifying the position of the judgements, meaning explanations and proofs on the one hand, and the propositions, definitions and verifications on the other hand.

3.1. The unit type. First, we introduce two canonical forms with trivial reduction rules:

$$\text{(Canonical)} \quad \text{unit} \Downarrow \text{unit} \quad \bullet \Downarrow \bullet$$

Next, we intend to make the judgement `unit type` evident; and this is done by defining what counts as a canonical verification of `unit` and when two such verifications are equal. To this end, we say that \bullet is a canonical verification of `unit`, and that it is equal to itself. I wish to emphasize that this is the entire definition of the type: we have introduced syntax, and we have defined the canonical forms, and there is nothing more to be done.

In the presentations of type theory which are currently in vogue, a type is “defined” by writing out a bunch of inference rules, but in type theory, the definitions that we have given above are prior to the rules, which are justified in respect of the definitions and the meaning explanations of the judgements. For instance, based on the meaning of the various forms of sequent judgement, the following rule schemes

are justified:

$$\begin{array}{c} \overline{\Gamma \gg \text{unit type}} \quad \overline{\Gamma \gg \text{unit} = \text{unit type}} \\ \overline{\Gamma \gg \bullet \in \text{unit}} \quad \overline{\Gamma \gg \bullet = \bullet \in \text{unit}} \end{array}$$

Each of the assertions above has evidence of a certain kind; since the justification of these rules with respect to the definitions of the logical constants and the meaning explanations of the judgements is largely self-evident, we omit it in nearly all cases. It is just important to remember that it is not the rules which define the types; a type A is defined in the course of causing the judgement $A \text{ type}$ to become evident. These rules merely codify standard patterns of use, nothing more, and they must each be justified.

3.2. The empty type. The empty type is similarly easy to define. First, we introduce a constant:

$$\text{(Canonical)} \quad \text{void} \Downarrow \text{void}$$

To make the judgement void type evident, we will say that there are no canonical verifications of void , and be done with it. This definition validates some further rules schemes:

$$\begin{array}{c} \overline{\Gamma \gg \text{void type}} \quad \overline{\Gamma \gg \text{void} = \text{void type}} \\ \frac{\Gamma \gg M \in \text{void}}{\Gamma \gg \mathcal{J}} \end{array}$$

The last rule simply says that if we have a verification of void , then we may conclude any judgement whatsoever. Remember that the inference rules are just notation for an *evident* hypothetical judgement, e.g. $\Gamma \gg \mathcal{J} (\Gamma \gg M \in \text{void})$.

Note that we did not introduce any special constant into the computation system to represent the elimination of a verification of void (in proof-theoretic type theories, this non-canonical form is usually called $\text{abort}(R)$). This is because, computationally speaking, there is never any chance that we should ever have use for such a term, since we need only consider the evaluation of closed terms (which is guaranteed by the meaning explanations), and by its very definition, there can never be a closed verification of void .

3.3. The cartesian product of a family of types. This will be our first dependent type, and it will likewise be the first example of a non-trivial addition to the computation system. First, let us add our canonical and non-canonical forms

and their reduction rules:

$$\begin{array}{l}
 \text{(Canonical)} \quad \Pi(A; x.B) \Downarrow \Pi(A; x.B) \quad \lambda(x.E) \Downarrow \lambda(x.E) \\
 \text{(Non-canonical)} \quad \frac{M \Downarrow \lambda(x.E) \quad [N/x]E \Downarrow M'}{\text{ap}(M; N) \Downarrow M'}
 \end{array}$$

This is the first time in this chapter that we have introduced a term former with binding structure; it should be noted that the intensional equality of expressions is up to alpha equivalence, and we will not pay attention to issues of variable renaming in our presentation.

We will make evident the following judgement scheme:

$$\frac{A \text{ type} \quad x : A \gg B \text{ type}}{\Pi(A; x.B) \text{ type}}$$

Or, written as a hypothetical judgement:

$$\Pi(A; x.B) \text{ type} \quad (A \text{ type}, x : A \gg B \text{ type})$$

This is to say, under the stated assumptions, we know what counts as a canonical verification of $\Pi(A; x.B)$ and when two such verifications are equal. We will say that $\lambda(x.E)$ is a canonical verification of $\Pi(A; x.B)$ just when we know $x : A \gg E \in B$; moreover, that two verifications $\lambda(x.E)$ and $\lambda(y.E')$ are equal just when $z \in A \gg [z/x]E = [z/y]E' \in [z/x]B$.

By the meaning of the sequent judgement, this is to say that a lambda expression must be functional with respect to its domain (i.e. it must take equals to equals). We did not need to hypothesize directly two elements of the domain and their equality because this is part of the meaning explanation for the sequent judgement already. Likewise, two lambda expressions are equal when equal inputs yield equal results in both.

The familiar inference rules, which codify the standard mode of use for the family cartesian product, are justified by this definition:

$$\begin{array}{l}
 \frac{\Gamma \gg A \text{ type} \quad \Gamma, x : A \gg B \text{ type}}{\Gamma \gg \Pi(A; x.B) \text{ type}} \quad \frac{\Gamma \gg A = A' \text{ type} \quad \Gamma, z : A \gg [z/x]B = [z/y]B' \text{ type}}{\Gamma \gg \Pi(A; x.B) = \Pi(A'; y.B') \text{ type}} \\
 \frac{\Gamma, x : A \gg E \in B}{\Gamma \gg \lambda(x.E) \in \Pi(A; x.B)} \quad \frac{\Gamma, z : A \gg [z/x]E = [z/y]E' \in [z/x]B}{\Gamma \gg \lambda(x.E) = \lambda(y.E') \in \Pi(A; x.B)} \\
 \frac{\Gamma \gg M \in \Pi(A; x.B) \quad \Gamma \gg N \in A}{\Gamma \gg \text{ap}(M; N) \in [N/x]B} \quad \frac{\Gamma \gg M = M' \in \Pi(A; x.B) \quad \Gamma \gg N = N' \in A'}{\Gamma \gg \text{ap}(M; N) = \text{ap}(M'; N') \in [N/x]B}
 \end{array}$$

Note that the type equality rule scheme that we gave above is structural; it is validated by the meaning explanations, but it is by no means the full totality of

possible equalities between family cartesian product types, which is extensional in this theory.

It will be instructive to explicitly justify the application rule above with respect to the meaning explanations, since I have claimed that such rules are posterior to the definitions we expounded prior to giving these rules.

PROOF. It suffices to consider the case that $\Gamma \Downarrow \cdot$, because hypotheses may always be added to the context (this is called weakening). And so, by the meaning of the sequent judgement at the empty context, the rule amounts to the assertion

$$\mathbf{ap}(M; N) \in [N/x]B \quad (M \in \Pi(A; x.B), N \in A)$$

By the meaning explanation for hypothetical judgement, and the definition of the family cartesian product type, we know that $M \Downarrow \lambda(x.E)$ for some E such that we know $x : A \gg E \in B$; from the meaning of the sequent judgement and the notion of an environment for a context, we can conclude $|_L [L/x]E \in [L/x]B$ ($L \in A$). On the other hand, from the computation rules, we know that if for some particular E' , $[N/x]E \Downarrow E'$, then we know $\mathbf{ap}(\lambda(x.E); N) \Downarrow E'$; to demonstrate the evidence of the premise, we may instantiate L at N to know $[N/x]E \in [N/x]B$, whence by the meaning of membership, we know that there exists some canonical E' such that $[N/x] \Downarrow E'$.

□

Proof Theories and Logical Frameworks

1. Proof-theoretic type theory

1.1. Analytic and synthetic judgement. A synthetic judgement is one for which the experience of *coming to know it* necessarily entails some knowledge which is not implicit in the statement of the judgement; on the other hand, to know an *analytic* judgement is to know it purely on the basis of the information contained inside it. So analytic judgements are decidable, since if they may become evident, it will be purely on the basis of their own content; whereas synthetic judgements become evident to someone when they have obtained some particular evidence for them.

A logical theory has, then, both analytic and synthetic judgements; the judgement $P \text{ prop}$ is analytic, since its evidence follows from the definition of P , whereas the assertion of $P \text{ true}$ entails the knowledge of some extra information, namely a verification of P . When we have extended the logical theory to a type theory in the manner of the previous chapter, the judgement $M \in P$ is also synthetic, since $M \in P$ is not self-evident in general.

But why is it not enough to assert that M verifies P to know whether $M \in P$? It suffices to define a P such that one cannot decide in general whether some term is a verification of it. Let us define the propositional symbol P , and we intend to know the judgement $P \text{ prop}$, whose meaning is to be expanded as follows:

To know $P \text{ prop}$ is to know counts as a canonical verification of P .

We will say, then, that \bullet is a canonical verification of P just when Goldbach's conjecture is true. Then it comes immediately that the judgement $M \in P$ may not be known or refuted on its own basis, nor even the judgement $\bullet \in P$, since they depend on a proposition whose truth is not known:

To know $M \in P$ is to know that $M \Downarrow M'$ to a canonical verification of P .
 \rightsquigarrow To know $M \in P$ is to know that $M \Downarrow \bullet$ such that Goldbach's conjecture is true.

1.2. Proof of a judgement vs. verification of a proposition. Because the judgement $M \in P$ is synthetic, we cannot say that it gives rise to a proof

theory for the logic, since the core judgement of a proof theory $M : A$ must be analytic, in order to avoid the infinite regress of a proof theory requiring a proof theory requiring a proof theory, and so on.

The notion of verification of a proposition could never be the same as proof anyway, except in the most trivial circumstances, since a verification is meant to be an effective operation which realizes the truth of a proposition, and no constraints whatsoever (termination, totality, etc.) are placed on these operations except those which come from the meaning of the judgements (see [4], [16], [17]).

So a proof theory is necessarily intensional, and its judgements are to be analytic/decidable. What is it, then, that we have considered so far which corresponds with a proof M such that $M : P$ in a proof theory? As discussed above, M is not merely a term such that $M \in P$, since this is not in general enough information to know whether M is a proof. In fact, M must comprise all the logical inferences which led to the knowledge that P is true, and so a meaning explanation for the judgement $M : P$ in a proof theory immediately suggests itself:

To know $M : P$ is to know that M is evidence (demonstration, proof, derivation) of the judgement P true.

And so the term domain of the proof theory is not the same as the one that we have considered so far; it must consist in terms which represent traces of the inferences made in the course of knowing the judgements of a logical theory. There is a sense in which one can consider the types of a proof theory to interpret the judgements of the logical theory, and this methodology is called “judgements as types” (and this implies “derivations as terms”).

What I am calling a “proof-theoretic type theory” is a type theory of the sort used in the proof assistants Agda, Coq and Idris, whereas the kind of type theory that I have described in the previous sections, the one based on meaning explanations, underlies the proof assistant Nuprl.

The proof-theoretic type theories on the one hand are often called “intensional” and the computational type theories on the other hand are usually “extensional”; these characterizations are certainly true, though they are not *essential*; moreover, I fear that comparing one of the former with one of the latter is not quite fair, since there is not any clear analogy to be had. That is to say, the judgement $M \in P$ is a judgement which is added to a logical theory and its meaning is (briefly) “ M evaluates to a canonical verification of P ”, whereas $M : P$ cannot be construed as a judgement added to a logical theory. Instead, it must be understood as part of a (proof) theory which is overlayed atop an existing logical theory; it is possible to understand the theory which contains the judgement $M : P$ to be a metatheory, or

logical framework, for the theory which contains the judgement P *true*, which can be construed as the “object language”.

In short, the judgements $M \in P$ and $M : P$ are unrelated to each other in two respects: firstly, that they have different meanings, and secondly that the one is at the same level as the judgements of a logical theory, whereas the latter is a judgement in a theory which is defined over a logical theory.

2. Martin-Löf's equational logical framework

To make this more concrete, let us expound a proof theoretic type theory called **MLLF**, which stands for “Martin-Löf's (equational) logical framework”;¹ in the course of introducing each type, we will specify which judgement of the underlying logical theory it is meant to interpret.

We start with four categorical judgements:

Judgement Form	Pronunciation
$\alpha : \textit{type}$	α is a type
$\alpha = \beta : \textit{type}$	α and β are equal types
$M : \alpha$	M is of type α
$M = N : \alpha$	M and N are equal at type α

But we have not defined the meaning of the judgements; let us do so below:

To know $\alpha : \textit{type}$ is to know what counts as an object of type α , and when two such objects are equal.

For now, we'll leave the question of what is an “object” as abstract; in many cases, types will represent judgements of a logical theory, and the objects will be the derivations (demonstrations, proofs) of those judgements.

To know $\alpha = \beta : \textit{type}$ is to know that any object of type α is also an object of type β , and two equal objects of type α are equal as objects of type β (necessarily presupposing $\alpha : \textit{type}$ and $\beta : \textit{type}$).

To know $M : \alpha$ is to know that M is an object of type α (necessarily presupposing $\alpha : \textit{type}$).

To know $M = N : \alpha$ is to know that M and N are equal objects of type α (necessarily presupposing $M : \alpha$ and $N : \alpha$).

In addition to the above judgements, we will need contexts (with their well-formedness judgement $\Gamma \textit{ctx}$) and an intensional sequent judgement $\boxed{\Gamma \vdash \mathcal{J}}$; their

¹For a detailed overview of Martin-Löf's equational logical framework, see [15].

meanings here will differ from the sequent judgements of computational type theory, in that they must mean proof-theoretic derivability, rather than semantic consequence.

At this point, we may begin adding types to the logical framework. In practice, most types which we will introduce in the logical framework will be defined in terms of a judgement of the logical theory which lies below it. For instance, hypothetical judgement in the logical theory is represented by a function type in the logical framework, $(x : \alpha)\beta$, whose typehood is meant to be evident under the following circumstances

$$\frac{\alpha : \text{type} \quad x : \alpha \vdash \beta : \text{type}}{(x : \alpha)\beta : \text{type}}$$

Or as a hypothetical judgement, $(x : \alpha)\beta : \text{type} \ (\alpha : \text{type}, x : \alpha \vdash \beta : \text{type})$.

Now, to know this judgement is to know that under the circumstances we know what is an object of type α and when two such objects are equal, and that if we have such an object x of type α , we know what an object of type β is, and when two such objects are equal—then we know what an object of type $(x : \alpha)\beta$ is, and moreover, for any two objects y, z of type α , that $[y/x]\beta$ and $[z/x]\beta$ are equal as types. To make this evident, then, we will say that under those circumstances an object of type $(x : \alpha)\beta$ is an object $[x]M$ such that one knows $x : \alpha \vdash M : \beta$ and $|_{x,y} [y/x]M = [z/x]M : [y/x]\beta \ (y = z : \alpha)$; furthermore, two such objects are equal just when they yield equal outputs for equal inputs.

Then, for each atomic proposition P , we can easily define a type $\text{Prf}(P)$, as follows. Under the circumstances that $P \text{ prop}$ in the logical theory, then $\text{Prf}(P) : \text{type}$ in the logical framework, since we will define an object of type $\text{Prf}(P)$ to be a derivation of $P \text{ true}$; beyond reflexivity, further definitional equalities can be added to reflect the harmony of introduction and elimination rules.

Now, the definitions we have given for the types above are “intuitively” correct, but they actually fail to satisfy the meaning explanation that we have given for $\alpha : \text{type}$, because they do not take into account neutral terms. In the following sections, we will investigate this problem in more detail and propose a solution.

2.1. What is an “object”? It is time to revisit what it means to be an “object” of a type in the proof-theoretic type theory; we must note how this will necessarily differ from what it meant to be a “verification” of a proposition in the previous sections. Namely, a verification of a proposition is either a *canonical verification* of that proposition (and what sort of thing this might be is known from the presupposition $P \text{ prop}$), or it is a means of getting such a canonical verification (i.e. a term which evaluates to a canonical verification).

On the other hand, what we have called an “object” of type P is quite different, since in addition to the possibility that it is a canonical proof of the judgement

P true, it may also be *neutral* (i.e. blocked by a variable); we will call this “normal” rather than “canonical”. Why does this happen?

In order to keep the judgement $M : A$ analytic (decidable), its meaning explanation can no longer be based on the idea of the computation of closed terms to canonical form; instead, we will consider the computation of open terms (i.e. terms with free variables) to *normal* form. The desire for $M : A$ to be analytic follows from our intention that it characterize a *proof theory*: we must be able to recognize a proof when we see one. But why are closed-term-based meaning explanations incompatible with this goal? Consider briefly the following judgement:

$$|_n M(n) \in P \ (n \in \mathbb{N})$$

To know this judgement is to know that $M(n)$ computes to a canonical verification of P whenever n is a natural number; when P 's use of n is not trivial, this amounts to testing an infinite domain (all of the natural numbers), probably by means of mathematical induction. The judgement is then clearly synthetic: to know it is, briefly, to have come up with an (inductive) argument that $M(N)$ computes to a canonical verification of P at each natural number n .

On the other hand, the judgement $n : \mathbb{N} \vdash M(n) : P$ must have a different meaning, one which admits its evidence or refutation purely on syntactic/analytic grounds. In essence, it is to know that $M(n)$ is a proof of P for any *arbitrary* object/expression n such that $n : \mathbb{N}$ (i.e., the only thing we know about n is that it is of type \mathbb{N} ; we do not necessarily know that it is a numeral).

3. A critique of MLLF

The type theory which we constructed in the previous section is to be considered a proof theory for a logic with the judgements P prop, P true and \mathcal{J} (\mathcal{J}'). There are a few reasons to be dissatisfied with this state of affairs, which I shall enumerate in this section.

3.1. Lack of computational content. Unlike the type theory in the first chapter, there is no built-in computational content. In a computational type theory which is defined by the verificationist meaning explanations, the computational content of terms is understood immediately by means of the $M \Downarrow M'$ relation; that is, computation is prior to the main judgements because their meaning explanations are defined in terms of evaluation to canonical form.

On the other hand, in the type theory above we did not give a primitive reduction relation; instead, we simply permitted the endowment of proofs with definitional equalities which reflect the harmony of introduction and elimination rules. That is, if we have known the judgement P true by means of an indirect argument (derivation), it must be the case that this derivation corresponds to a direct

one; we reflect this in the proof theory by defining the indirect derivation to be definitionally equal to the direct one.

However, this does not amount to computational content being present in terms: only *post facto* may the definitional equality be construed as giving rise to computation, through a metamathematical argument which shows that the definitional equality is confluent and can be used to define a functional normalization relation.

And this is the reason for the peculiarity of the proof-theoretic meaning explanations, namely that they do not include phrases like “evaluates to a canonical...”, since evaluation may only be understood after taking the meanings of the judgements ($\alpha : type$, $\alpha = \beta : type$, $M : \alpha$, $M = N : \alpha$) as giving rise to a closed formal system which is susceptible to metamathematical argument: to refer to evaluation in the meaning explanations for the core judgements, then, would be impredicative.

3.2. Modularity of definition. By the same token, the distinction between canonical (direct) and non-canonical (indirect) proof may not be understood as a core notion in the theory, but must be understood separately, secondarily. Why is this a problem? It means that the definition of each type must be made with the full knowledge of the definitions of every other type; in essence, the open-ended nature of type theory is obliterated and one is forced into a fixed formal system; this is in addition to the fact that it causes the epistemic content of $\alpha : type$ for any type α to be extremely complicated.

To illustrate, let us consider as an example a type theory which has four type-formers: trivial truth \top , trivial falsity \perp , implication $(\alpha)\beta$, and conjunction $\alpha \& \beta$; we will then introduce the following terms to represent proofs: the trivial element \bullet , *reductio ad absurdum* $\text{abort}(\alpha; E)$, abstraction $[x : \alpha]E$, application $E(E')$, pairing $\langle E, E' \rangle$, and projections $\text{fst}(E)$, $\text{snd}(E)$.

If we will try to make the judgement $\top : type$ evident, the deficiencies of the formulation will immediately present themselves.

To know $\top : type$ is to know what counts as an object of type \top , and when two such objects are equal. An object of type \top , then, is either the expression \bullet , or an expression $\text{abort}(\top; E)$ such that we know $E : \perp$, or an expression $E(E')$ such that we know $E : (\alpha)\top$ and $E' : \alpha$, or an expression $\text{fst}(E)$ such that we know $E : \top \& \beta$ for some β , or an expression $\text{snd}(E)$ such that we know $E : \alpha \& \top$ for some α ; and we additionally have that \bullet is equal to \bullet , and ...

To save space, we elide the rest of the definition of equality for \top ; what we have seen so far already suffices to bring to light a serious problem: the definition of any type requires knowledge of the entire syntax of the theory. The judgement $\alpha : type$

may never be made evident in isolation, but must be done with full understanding of all the other types and their definitions.

Furthermore, to extend an existing theory with a new type, the definitions of every other type must be rewritten to account for the elimination forms of the new type.

4. A way forward: verifications & uses

The second critique of **MLLF** may be partially addressed by fragmenting type theory into a logic of *verifications & uses*: instead of a type being defined by its introduction rules, it must be simultaneously defined by its introduction rules (verifications) and its elimination rules (uses). In practice, this amounts to a standard technique known as *bidirectional type checking*.

The semantic priority of the forms of judgement also changes drastically: the sequent judgement must in this case be explained *before* the categorical judgements; moreover, sequents may no longer be explained modularly in terms of general and hypothetical judgement, since the latter amounts to *semantic* consequence (admissibility), whereas the meaning of a sequent in a proof theory should be *syntactic* consequence (derivability).

Because the target theory lacks computation, it is necessary to rule out redexes from terms syntactically, but this complicates the definition of substitution; to address this, Watkins introduced in [18] a technique known as *hereditary substitution*, which is a family of syntax-directed (algorithmic) judgements which contract redexes along the way, guaranteeing canonical form in their outputs. Both bidirectional type checking and hereditary substitutions have been used to great effect in the descendants of the Edinburgh Logical Framework (see [10]).

The first critique, the lack of computational content, is more difficult to address. Roughly, the right way to do it is to replace the notion of the evaluation of closed terms to canonical form with the evaluation of *open* terms to normal form. Peter Dybjer demonstrates in [5] how this technique may be used to endow the Calculus of Constructions with a meaning explanation, albeit necessarily of a very different kind than we have considered here.

Bibliography

- [1] *Intuitionistic Type Theory*. Bibliopolis, 1984.
- [2] Stuart F. Allen, Mark Bickford, Robert L. Constable, Richard Eaton, Christoph Kreitz, Lori Lorigo, and E. Moran. Innovations in computational type theory using nuprl. *J. Applied Logic*, 4(4):428–469, 2006.
- [3] Robert L. Constable, Stuart F. Allen, Mark Bromley, Rance Cleaveland, J. F. Cremer, R. W. Harper, Douglas J. Howe, Todd B. Knoblock, N. P. Mendler, Prakash Panangaden, James T. Sasaki, and Scott F. Smith. *Implementing mathematics with the Nuprl proof development system*. Prentice Hall, 1986.
- [4] Michael Dummett. *Elements of intuitionism*, volume 39 of *Oxford Logic Guides*. The Clarendon Press Oxford University Press, New York, second edition, 2000.
- [5] Peter Dybjer. Program testing and the meaning explanations of intuitionistic type theory. In Peter Dybjer, Sten Lindström, Erik Palmgren, and Göran Sundholm, editors, *Epistemology versus Ontology*, volume 27 of *Logic, Epistemology, and the Unity of Science*, pages 215–241. Springer, 2012.
- [6] Johan Georg Granström. *Treatise on Intuitionistic Type Theory*. Springer Publishing Company, Incorporated, 2013.
- [7] Robert Harper. Constructing type systems over an operational semantics. *J. Symb. Comput.*, 14(1):71–84, 1992.
- [8] Robert Harper. *Practical Foundations for Programming Languages*. Cambridge University Press, New York, NY, USA, 2012.
- [9] Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *J. ACM*, 40(1):143–184, January 1993.
- [10] Robert Harper and Daniel R. Licata. Mechanizing metatheory in a logical framework. *J. Funct. Program.*, 17(4-5):613–673, July 2007.
- [11] Per Martin-Löf. Constructive mathematics and computer programming. In L. Jonathan Cohen, Jerzy Łoś, Helmut Pfeiffer, and Klaus-Peter Podewski, editors, *Logic, Methodology and Philosophy of Science VI, Proceedings of the Sixth International Congress of Logic, Methodology and Philosophy of Science, Hannover 1979*, volume 104 of *Studies in Logic and the Foundations of Mathematics*, pages 153–175. North-Holland, 1982.
- [12] Per Martin-Löf. Truth of a proposition, evidence of a judgement, validity of a proof. *Synthese*, 73(3):407–420, 1987.
- [13] Per Martin-Löf. Analytic and synthetic judgements in type theory. In Paolo Parrini, editor, *Kant and Contemporary Epistemology*, volume 54 of *The University of Western Ontario Series in Philosophy of Science*, pages 87–99. Springer Netherlands, 1994.
- [14] Per Martin-Löf. On the meanings of the logical constants and the justifications of the logical laws. *Nordic Journal of Philosophical Logic*, 1(1):11–60, 1996.
- [15] B. Nordström, K. Petersson, and J.M. Smith. *Programming in Martin Löf’s Type Theory*, volume 7 of *Monographs on Computer Science*. Oxford University Press, 1990.

- [16] Dag Prawitz. Truth and proof in intuitionism. In Peter Dybjer, Sten Lindström, Erik Palmgren, and Göran Sundholm, editors, *Epistemology versus Ontology*, volume 27 of *Logic, Epistemology, and the Unity of Science*, pages 45–67. Springer, 2012.
- [17] A.S. Troelstra and D. van Dalen. *Constructivism in Mathematics*. North-Holland, 1988. 2 volumes.
- [18] Kevin Watkins, Iliano Cervesato, Frank Pfenning, and David Walker. A Concurrent Logical Framework: The Propositional Fragment. *Types for Proofs and Programs*, pages 355–377, 2004.