

McTT: A Verified Kernel for a Proof Assistant

JUNYOUNG JANG, McGill University, Canada

ANTOINE GAULIN, McGill University, Canada

JASON Z. S. HU*, Amazon, USA

BRIGITTE PIENTKA, McGill University, Canada

Proof assistants based on type theories have been widely successful from verifying safety-critical software to establishing a new standard of rigour by formalizing mathematics. But these proof assistants and even their type-checking kernels are also complex pieces of software, and software invariably has bugs, so why should we trust such proof assistants?

In this paper, we describe the McTT (Mechanized Type Theory) infrastructure to build a verified implementation of a kernel for a core Martin-Löf type theory (MLTT). McTT is implemented in Rocq and consists of two main components: In the theoretical component, we specify the type theory and prove theorems such as normalization, consistency and injectivity of type constructors of MLTT using an untyped domain model. In the algorithmic component, we relate the declarative specification of typing and the model of normalization in the theoretical component with a functional implementation within Rocq. From this algorithmic component, we extract an OCaml implementation and couple it with a front-end parser for execution. This extracted OCaml code is comparable to what a skilled human programmer would have written and we have successfully used it to type-check a series of small-scale examples.

McTT provides a fully verified kernel for a core MLTT with a full cumulative universe hierarchy. Every step in the compilation pipeline is verified except for the lexer and pretty-printer. As a result, McTT serves both as a framework to explore the meta-theory of advanced type theories and to investigate optimizations of and extensions to the type-checking kernel.

CCS Concepts: • **Theory of computation** → **Type theory**; *Logic and verification*; **Program verification**; *Denotational semantics*; **Constructive mathematics**.

Additional Key Words and Phrases: Type Theory, Verified Implementation, Martin-Löf Type Theory, Rocq, Normalization-by-Evaluation

ACM Reference Format:

Junyoung Jang, Antoine Gaulin, Jason Z. S. Hu, and Brigitte Pientka. 2025. McTT: A Verified Kernel for a Proof Assistant. *Proc. ACM Program. Lang.* 9, ICFP, Article 242 (August 2025), 32 pages. <https://doi.org/10.1145/3747511>

1 Introduction

Over the past decades, proof assistants based on type theories have been widely used to develop safety-critical software (e.g. COMPCERT [Leroy 2009a,b]). They have also been used to formalize a significant portion of mathematics with the goal of establishing a new standard of rigour [Avigad and Harrison 2014] in this field. For example, The Mathlib Community [2020] in the LEAN proof assistant has roughly half a million lines of code and contains formalizations of many non-trivial mathematics, from number theory to perfectoid spaces [Buzzard et al. 2020]. These endeavors are

*Jason Z. S. Hu made his contribution during his Ph.D. study at McGill University, Canada.

Authors' Contact Information: Junyoung Jang, junyoung.jang@mail.mcgill.ca, School of Computer Science, McGill University, Montréal, Québec, Canada; Antoine Gaulin, antoine.gaulin@mail.mcgill.ca, School of Computer Science, McGill University, Montréal, Québec, Canada; Jason Z. S. Hu, zhonhu@amazon.com, Amazon, Seattle, Washington, USA; Brigitte Pientka, bpientka@cs.mcgill.ca, School of Computer Science, McGill University, Montréal, Québec, Canada.

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2475-1421/2025/8-ART242

<https://doi.org/10.1145/3747511>

testimony to the success of proof assistants in practice. Nevertheless, proof assistants are complex and unavoidably have bugs, so why should we trust a proof assistant and its certification of a proof?

To minimize such concerns, proof assistants are typically designed to rely on a small, trusted kernel to verify a proof. This design approach narrows our focus to this trusted core, which is supposed to be close to its well-understood, underlying theory, e.g., Calculus of Inductive Construction (CIC) [Coquand and Paulin 1988; Paulin-Mohring 1993] or Martin-Löf type theory (MLTT) [Martin-Löf 1975]. However, in practice, kernels of modern proof assistants are no longer small, spanning over tens of thousands of lines of code. RocQ’s kernel spans approximately 41K lines of OCaml code, while LEAN’s kernel consists of approximately 8K lines of C++ code. Furthermore, the kernel might significantly deviate from and extend the core type theories in non-trivial ways, both to experiment with new features and to support more practical features such as meta-variables for partial proofs or asynchronous type-checking. As a result, all type-theoretic proof assistants have encountered major bugs: In RocQ, on average, one critical bug has been found every year. LEAN has experienced both soundness bugs¹ and segmentation fault.²

In this work, we investigate a long-standing question: *How can we build a verified type-checker for type-theoretic proof assistants?* – Two fundamental aspects need to be answered:

1) *Check the Theory.* The type theory underlying a proof assistant is itself a full-scale, complex system. Hence, verifying the type-checking kernel should start with formalizing the type theory itself, together with its meta-theory. Such an endeavour not only increases the trust in the type-theoretic foundation itself, but also such puts great pressure on many aspects of a proof assistant – from stretching the boundaries of the theoretical foundations, to understanding the good practices of developing such meta-theoretic results, and testing its implementation [Chapman 2008].

The most frequently adapted proof technique to prove normalization of MLTT is Tait’s reducibility candidates [Tait 1967] (see, for example, [Abel et al. 2023, 2018; Adjedj et al. 2024; Pujet and Tabareau 2022, 2023]). These mechanizations often require significant technical set-up and time investment. Moreover, this style of normalization only gives us β normal forms.

An alternative method is normalization-by-evaluation (NbE), where we use a mathematical model that includes a notion of computation. Such NbE proofs have the advantage of directly giving a normalization algorithm [Berger and Schwichtenberg 1991] and are often much more compact. For example, Hu et al. [2023] mechanize an NbE algorithm for MLTT extended with a necessity modality based on Abel [2013]’s untyped domain model in just 11K lines of Agda code. In RocQ, Wieczorek and Biernacki [2018] lead the effort of mechanizing a variant of MLTT using NbE.

However, there is a subtle issue that arises in these mechanizations of type theories: existing proofs in Agda [Abel et al. 2018; Hu et al. 2023; Pujet and Tabareau 2023] critically rely on induction-recursion [Dybjer and Setzer 2003] to formalize the semantics. However, RocQ does not support induction-recursion. For this reason, prior RocQ mechanizations of MLTT (e.g. [Adjedj et al. 2024; Wieczorek and Biernacki 2018]) restrict MLTT to a finite number of universes. This has been a stumbling block in scaling mechanizations of MLTT in RocQ. This is unfortunate, since RocQ has more support for automation and better code extraction than Agda.

In addition, none of the previous mechanizations of normalization proofs for MLTT were developed with an eye towards extracting a practical type-checker and normalizer from its development. However, to establish a correspondence between a type-checker and its underlying theory, the declarative specification of the type theory must be developed hand in hand with an algorithmic one that forms the core of a usable type-checker.

¹<https://github.com/leanprover/lean4/issues/1433> and <https://github.com/leanprover/lean4/issues/2125>

²<https://github.com/leanprover/lean4/issues/5188>

2) *Check the Checker*. A verified type-checker for a (core) type theory provides the highest level of assurance and eliminates errors. The first notable verified kernel of a significant fragment of CIC goes back to [Barras and Werner \[1997\]](#). [Boespflug and Burel \[2012\]](#) and [Assaf et al. \[2016\]](#) propose to translate RocQ proofs into the logical framework $\lambda\Pi$ -Modulo, and then use that framework to certify the RocQ proofs. However, since the type-checker for $\lambda\Pi$ -Modulo is smaller than that for RocQ, we cannot establish meta-theoretic results about encodings within the framework.

An alternative approach is pursued in the MetaRocQ project, which aims to formalize RocQ in RocQ and provide tools for manipulating RocQ terms and developing certified plugins [\[Sozeau et al. 2020a,b, 2019\]](#). The main idea is to reflect the representation of RocQ objects in the implementation of RocQ into RocQ itself and to prove meta-theoretic properties such as subject reduction, canonicity, or weak call-by-value standardization. However, while there is a normalization proof for the CIC, there is no normalization proof for the system that is actually modelled in RocQ.

While one might argue that we cannot prove and hence mechanize the consistency of the entire RocQ's type theory in RocQ, one could still hope to mechanize a significant fragment of it.

Contribution. In this paper, we develop McTT (Mechanized Type Theory), a verified type-checker of MLTT in RocQ. Channeling the spirit of the CompCert project [\[Leroy 2009a,b\]](#), we start by formalizing the type theory together with its meta-theory (*theoretical component*). To extract a usable type-checker, we also implement a type-checker for our type theory in RocQ (*algorithmic component*) and prove it correct with respect to the specification in the theoretical component.

In *the theoretical component*, we *check the theory*. In particular, we formalize a version of MLTT that has natural numbers, Π types, a full cumulative universe hierarchy, and universe subtyping. This formalization includes the mechanized proof of logical consistency using normalization-by-evaluation (NbE), closely following the proof of [Abel \[2013\]](#). Compared to a similar previous mechanization effort [\[Wieczorek and Biernacki 2018\]](#) in RocQ, McTT formalizes a richer type theory. A key obstacle when mechanizing the meta-theory of MLTT with a full universe hierarchy is that its semantics is prevalently defined by using induction-recursion. We provide a principled way to formalize such definitions in RocQ by exploiting impredicativity. As such, the theoretical component of McTT not only increases our confidence in the core MLTT, but provides insights and derives good practices on how to formalize the meta-theory of type theories in proof assistants such as RocQ. Moreover, we develop the meta-theory of MLTT hand-in-hand with its algorithmic, functional implementation. This leads us to incorporate universe subtyping into MLTT. This is an extension which is usually supported in CIC, but not traditionally in MLTT. Universe subtyping is necessary to establish soundness and completeness of the algorithmic type checker.

In *the algorithmic component* of McTT, we *check the checker*. Specifically, we relate the model of normalization in the theoretical component and the actual implementations of a type checker and normalization engine, in order to prove that they are equivalent. This allows us to extract a fully verified type-checker to OCaml. The extracted code does not contain any proof witnesses and is comparable to what a skilled human programmer would have written. Furthermore, we connect the extracted type-checker with a front-end parser, so that we can read and type-check MLTT programs from source files using our certified type-checking kernel. In fact, we have used the McTT's extracted type-checker to type-check several small-scale examples to test that it works as intended. At this point, the only unverified component in the compiler pipeline is the lexer, as `ocamllex` does not support extraction to RocQ.

In summary, McTT provides a framework for building verified kernels for dependent type theories and bridges the gap between the underlying theory and its implementation. Concretely, it makes the following contributions:

- *Check the theory.* We present a mechanized semantics for a core MLTT with natural numbers, Π types, a full cumulative universe hierarchy, and covariant universe subtyping. In particular, we use an impredicative encoding with subtyping in Rocq's Prop universe to avoid induction-recursion. Using this encoding, we prove the completeness and soundness of the NbE algorithm using an untyped domain model [Abel 2013] and conclude the logical consistency of the type theory as a corollary.
- *Check the checker.* We implement an algorithmic type checker for our variant of MLTT in Rocq and prove that it is correct with respect to the theoretical specification. In particular, we prove soundness and completeness: an MLTT program is well-formed if and only if it is accepted by the type-checker.
- *Extract a verified type-checker.* From the algorithmic implementation of the type-checker, we extract a readable OCaml implementation which we have used to type-check several small-scale examples. We further embed this verified type-checker within our front-to-back pipeline where users are able to write a program in MLTT in a source file and then use our verified type-checker to check it.

The McTT infrastructure provides a platform for formalizing future extensions of MLTT and derive verified proof checkers. Furthermore, McTT allows us to experiment with new features even before we extend the theoretical and algorithmic components by directly modifying the extracted core type-checker written in OCaml.

The source code and a homepage are available on [Github](#).

2 The McTT Infrastructure

In this section, we give a high-level overview of McTT. A pictorial summary can be found in Fig. 1. The heart of the infrastructure is the *theoretical and algorithmic component* implemented in Rocq.

2.1 Theoretical Component in Rocq

As in prior work by Abel et al. [2018], we focus on a variant of MLTT that includes natural numbers with a primitive recursor and Π types. We model variables using de Bruijn indices and support explicit substitutions. This facilitates reasoning and is common in normalization proofs. While this core fragment of MLTT is well understood, we also push our understanding of how to mechanize the meta-theory of MLTT in Rocq. In particular, we make the following three contributions:

Impredicativity Instead of Induction-Recursion. Modeling a full cumulative universe hierarchy is intuitive: each universe is indexed by a natural number, and lower universes are contained in higher universes. However, mechanizing its meta-theory is non-trivial. Previous mechanizations of MLTT in Agda [Abel et al. 2018; Hu et al. 2023; Pujet and Tabareau 2023] rely on induction-recursion to formalize the semantics, but Rocq does not support induction-recursion. One common workaround to this issue, the Bove-Capretta method [Bove 2009; Bove and Capretta 2005], cannot be not directly applied to full universe hierarchy. Prior Rocq mechanizations of MLTT (e.g. [Adjedj et al. 2024; Wieczorek and Biernacki 2018]) in Rocq restrict MLTT to at most two universes to avoid this issue. A key contribution of our mechanization is an alternative, principled way of modeling the semantics using impredicativity instead of induction-recursion in Rocq. We combine the inductively defined semantics of types and the recursively defined semantics of terms into a single inductive definition of a *weak partial* functional relation whose output is unique up to logical equivalence. This relation itself captures the semantics of types and returns the semantics of terms. This encoding not only provides an alternative way to model semantics without induction-recursion, but also takes advantage of the fact that the output of the functional relation is unique up to logical equivalence to simplify proofs of theorems that are more complex with induction-recursion. This




Fig. 1. Structure of McTT

equivalence-based definition further avoids postulating propositional extensionality and reduces the use of functional extensionality (cf. Sec. 4.1).

Universe Subtyping: Pursuing a Complete Algorithmic Type-checking. Universe subtyping is usually supported in CIC, but not in much prior work in MLTT [Abel 2013; Coquand 2018; Hu et al. 2023]. Instead, MLTT simply states: if a type is well-formed at universe level l , then it is also well-formed at universe level $1 + l$. However, we need a more general notion of universe subtyping to establish the soundness and completeness of algorithmic type-checking w.r.t. the declarative typing. For example, suppose that we have the following expression:

$$(\lambda(x:\text{Type}@1).\lambda(y:\text{Type}@0).y) \text{Type}@0 : (\Pi(y:\text{Type}@0).\text{Type}@1)$$

We first introduce a function with two arguments, where x is not used, and apply the function to $\text{Type}@0$. Declaratively, this expression has type $\Pi(y:\text{Type}@0).\text{Type}@1$, because we apply cumulativity to the variable y of $\text{Type}@0$. On the other hand, in a type-checking algorithm, we first infer the type of the function and get its type $\Pi(x:\text{Type}@1).\Pi(y:\text{Type}@0).\text{Type}@0$. The function application leads to the overall type $\Pi(y:\text{Type}@0).\text{Type}@0$ for the expression, but it is not related to the desired type $(\Pi(y:\text{Type}@0).\text{Type}@1)$. In general, it is challenging for algorithmic type-checking to retrospectively insert cumulativity properly. Using universe subtyping, we can naturally resolve this issue: $\Pi(y:\text{Type}@0).\text{Type}@0$ is a subtype of $\Pi(y:\text{Type}@0).\text{Type}@1$. Thus, the program has the latter type.

For this reason, we include a *covariant* subtyping relation between universes that propagates only under the outputs of Π types. This extension is sufficient to prove the correctness of algorithmic typing, while keeping the theoretical development manageable.

This example illustrates the value of developing the theory hand in hand with algorithmic implementations of that theory. Algorithmic and implementation concerns often lead to (ad-hoc) changes in the underlying theory, however those changes are not always accounted for in the theoretical development. This can then potentially give rise to subtle soundness bugs that are hard to detect. Developing the algorithmic implementation in a tight loop with the theory within a proof assistant provides a strong net that ensures that the theory is in sync with its implementation.

A Compact Mechanization of NbE. A main meta-theoretical property that we prove about MLTT is normalization based on NbE. NbE offers important advantages compared to the common reducibility candidate approach. In particular, it results in a compact mechanization. For example, our theoretical component of McTT in Rocq is only 12K LoC. This is significantly smaller than the mechanization using reducibility candidates without a full universe hierarchy and without universe subtyping (26k LoC in Rocq) [Adjedj et al. 2024]. Furthermore, NbE immediately provides an evaluation algorithm for (well-typed) programs of the language. In Rocq, we can also extract a usable implementation of a type-checker in OCaml. Last, NbE returns $\beta\eta$ normal forms, which can be compared syntactically. This is in contrast to reduction/rewrite rules in the approaches based on reducibility candidates which only give β normal forms and still require a separate comparison algorithm to handle η expansions [Abel et al. 2018].

2.2 Algorithmic Component in Rocq

In the algorithmic component, we implement a type-checking algorithm. This component consists of three algorithms: a functional implementation of the NbE algorithm, a subtyping algorithm, and a type-checking algorithm. These algorithms are eventually extracted into OCaml and form the kernel of the McTT executable.

We adapt a syntax-directed type-checking algorithm to control the use of subtyping. An object has type A if we are able to infer a type B from it and B is a subtype of A . We hence replace checking whether the type B is convertible to A with checking the subtyping relation. To check the subtyping relation for the types A and B , we first normalize the types A and B using the NbE algorithm. We then check whether the universe levels in the normalized A are always smaller than the levels in the *covariant* positions of normalized B .

We prove the correctness of all the implementations in the algorithmic components with respect to the declarative formulations in the theoretical components. An important theorem of McTT is that an MLTT program is well-formed if and only if it is accepted by the type-checker, hence by the extracted proof checker. As a consequence, we can have full confidence in the extracted proof checker.

2.3 Integration into the Compilation Pipeline

Our theoretical and algorithmic components are connected to a front-end that generates an abstract syntax tree (AST) from a given source file. The OCaml driver of McTT begins with a lexer generated by `ocamllex`, which produces a stream of tokens from a source file. Since `ocamllex` does not support extraction to Rocq, the lexer remains an unverified component in the pipeline. The stream of tokens is then forwarded to the parser, which is Rocq code generated by parser generator `Menhir`. The parser converts the tokens into a concrete syntax tree (CST), if the source file is grammatically well-formed. The parser is proven correct and complete as `Menhir` also generates the corresponding theorems. Then, the elaborator converts the CST into an abstract syntax tree (AST). At the current stage, the elaborator only converts variables represented in strings into de Bruijn indices. In the future, the elaborator can perform other extensions, such as module name expansions. The

soundness property of the elaborator is captured by a well-scopedness theorem, which says that if all open variables in strings of a CST are well-scoped, then this CST must elaborate to an AST.

The AST is used by both the theoretical and algorithmic components; in particular, it is passed to the extracted type-checker. Using the same AST across all components ensures that every step in the pipeline, including normalization, subtyping, and type-checking, is correct by construction, supported by a sequence of theorems. Finally, a normal form of the input program is pretty-printed.

McTT is designed to minimize the use of unverified components. Currently, unverified components involve the boundaries between the formal and the informal, e.g. the lexer and the pretty-printer. All other components are entirely implemented in RocQ, where we establish a series of soundness and completeness theorems about the parser, the elaborator, and more importantly the type-checker. The infrastructure provided by McTT is a significant stepping stone towards building verified proof checkers for type-theoretic proof assistants.

3 A Core Martin-Löf Type Theory with Universe Hierarchy and Universe Subtyping

In this section, we define the syntax of MLTT, its semantics, and finally the NbE procedure.

3.1 Syntactic Definitions of MLTT

3.1.1 Syntax of MLTT. The syntax of MLTT is similar to that presented by Abel [2013]. It includes natural numbers with a primitive recursor, Π types, closures using explicit substitutions, and a full universe hierarchy.

Variables			x, y, r
De Bruijn indices	$\mathbb{N}$	$\ni$	d
Universe levels	$\mathbb{N}$	$\ni$	i, j
Contexts	Ctx	$\ni$	$\Gamma, \Delta ::= \cdot \mid \Gamma, x : A$
Expressions	Exp	$\ni$	$A, B, C, M, N ::= x_d \mid \text{Nat} \mid \text{Type}@i \mid \Pi(x:A).B$ $\mid \text{zero} \mid \text{succ } M \mid \text{rec}_{x:A} M_{\text{zero}} (y, r.M_{\text{succ}}) N$ $\mid \lambda(x:A).M \mid M N \mid M[\sigma]$
Substitutions	Subst	$\ni$	$\sigma, \delta ::= \text{id} \mid \text{wk} \mid \sigma \circ \delta \mid \sigma, M/x_0$

In the RocQ mechanization, we use de Bruijn indices to represent names. However, for clarity of presentation, we use abstract names in the text. When we need to refer to de Bruijn indices, we put a subscript to the variable, e.g. x_d . The same name is used to denote the same variable with potentially different de Bruijn indices to help readers distinguish variables more easily. We use M and N for expressions that are terms and A, B , and C for those that are types.

The syntax for natural numbers (Nat) includes zero and $\text{succ } M$, the successor of M . The recursion principle $\text{rec}_{x:A} M_{\text{zero}} (y, r.M_{\text{succ}}) N$ takes four parameters. First, N is the natural number being recursed on. This number is called *scrutinee*. The base case is handled by M_{zero} , when N is equivalent to zero. If N is $\text{succ } N'$, then the recursion is continued by $y, r.M_{\text{succ}}$. In this case, the open variable y is the predecessor, i.e. N' in this case, and r denotes the recursive call on N' . The type $x:A$ represents the *motive*, the return type of the recursion. The overall type of the recursion is $A[N/x]$, i.e. A with x substituted by N .

Next, we have standard dependent function spaces, $\Pi(x:A).B$. The introduction form $\lambda(x:A).M$ binds a new variable x (or x_0 in de Bruijn indices) of type A that can occur in the body M of the function, while the elimination form $M N$ applies the function M to the argument N . Since we have dependent types, the result type from function application is $B[N/x]$.

Explicit substitutions come into play via closures, denoted $M[\sigma]$, which represent the application of the substitution σ to the expression M . We define substitutions following standard practice




Fig. 2. Selected rules for MLTT

(see, for example, [Abel \[2013\]](#)) using four possible cases. The identity substitution id is a no-op to expressions, and is the left and right identity of composition $\sigma \circ \delta$. The weakening substitution wk weakens the topmost variable, i.e. that of de Bruijn index 0. Finally, we can also extend a substitution with an expression via $\sigma, M/x$.

3.1.2 Judgements of MLTT. Our variant of MLTT is given by seven mutually defined judgements:

- $\boxed{\vdash \Gamma}$: Γ is well-formed.
- $\boxed{\Gamma \vdash M : A}$: M is a well-formed expression of type A in context Γ .
- $\boxed{\Gamma \vdash \sigma : \Delta}$: σ is a well-formed substitution from Δ to Γ .
- $\boxed{\Gamma \vdash M \approx M' : A}$: M and M' are equivalent expressions of type A in context Γ .
- $\boxed{\Gamma \vdash \sigma \approx \sigma' : \Delta}$: σ and σ' are equivalent substitutions from Δ to Γ .
- $\boxed{\Gamma \vdash A \subseteq B}$: Type A is a subtype of B in context Γ .
- $\boxed{\vdash \Gamma \subseteq \Delta}$: Γ is a sub-context of Δ . This judgement generalizes the subtyping judgement.

The rules for MLTT with explicit substitutions are mostly standard (see Fig. 2). We only show the core fragment consisting of functions, function applications, Π -types, and closures. A more exhaustive definition of MLTT can be found in Appendix A (and [its mechanized definition](#)).

In the rules, the **highlighted** premises ensure that well-typed programs always have well-formed types. We prove that they are redundant with a presupposition lemma [[Harper and Pfenning 2005](#)], essentially stating that each subcomponent of a judgement must also be well-formed. Since each judgement has different subcomponents, each of them gets their own presupposition lemma. For example, the presupposition for the typing judgement states that if $\Gamma \vdash M : A$, then $\vdash \Gamma$ and $\Gamma \vdash A : \text{Type}@i$ for some i . For the rest of our discussion, we will ignore the highlighted premises.

Following RocQ's core type theory, we replace the standard type conversion rule with one that allows subsumption and we extend the standard MLTT with a covariant universe subtyping judgement. The subtyping judgement admits the subsumption of smaller universes in larger ones

and propagates this relation down to the output types of Π types (but not the input types). In this way, we implement a full *cumulative* universe hierarchy. The subtyping judgement allows a judgement that holds for a subtype to also hold for its supertype. For example, the following is the subsumption rule for typing:

$$\frac{\Gamma \vdash M : A \quad \Gamma \vdash A \subseteq A'}{\Gamma \vdash M : A'}$$

This rule says that if M has type A and A is a subtype of A' , then M also has type A' . In the standard MLTT, the associated rule only checks the equivalence between A and A' . In contrast to the symmetry of equivalence, subtyping is asymmetric or uni-directional. In other words, we can only go from A to A' but not the other way.

The reason for introducing subtyping to this system is grounded in establishing completeness of the type-checker w.r.t. the declarative typing rules. Were we to eliminate universe subtyping, the type-checker would admit more terms than the typing rules. Consider a weaker form of cumulativity employed by Abel [2013]; Hu et al. [2023], which does not involve subtyping and is what we originally formulated:

$$\frac{\Gamma \vdash A : \text{Type}@i}{\Gamma \vdash A : \text{Type}@1+i}$$

Unfortunately, it turns out to be extremely challenging to develop a *complete* type-checking algorithm w.r.t this rule. As the example in the introduction illustrates, negative types, such as Π -types, are sensitive to when we apply this subtyping rule.

The bottom of Fig. 2 defines the subtyping rules. There are only four rules. The first rule says that equivalence between expressions is subsumed by subtyping, and it is effectively a reflexivity rule. The second rule is transitivity, which is expected for a subtyping relation. The third rule admits a smaller universe to be a subtype of a larger universe, thus supporting cumulativity. The final rule allows propagation of subtyping in the output types of Π types. The input types A and A' are equivalent since we do not support contravariant subtyping at this point, following Rocq.

Finally, the subtyping judgement between types is generalized pointwise to context subtyping, denoted as $\vdash \Gamma \subseteq \Delta$. In this judgement, the types in Γ are pointwise subtypes of those in Δ .

3.2 Semantics of MLTT

Next, we consider the semantic model of MLTT, which is intuitively just the untyped λ -calculus. Specifically, we use an untyped domain model à la Abel [2013] to formulate the NbE algorithm. This NbE algorithm has two steps:

- (1) First, we evaluate a well-typed term into a domain value (Sec. 3.3.1). During this process, we reduce away the β redexes.
- (2) Second, we read from the domain value back to the syntax of an object in normal form. During this process, we perform a type-directed η expansion.

Thus, combining both steps, we obtain the $\beta\eta$ normal form of the given well-typed term. We can define syntactic β -short forms using a standard notion of neutral and normal forms:

Neutral forms $\text{Ne} \ni V ::= x_d \mid \text{rec}_{x.W} W_{\text{zero}}(y, r.W_{\text{succ}}) V \mid V W$

Normal forms $\text{Nf} \ni W ::= V \mid \text{Nat} \mid \text{Type}@i \mid \Pi(x:W_A).W_B \mid \text{zero} \mid \text{succ } W \mid \lambda(x:W_A).W_M$

Neutral forms are either variables or elimination principles blocked by neutral forms. Normal forms are either neutral, type constructors, or introduction principles applied to normal forms. Note that these normal forms are not necessarily η -long. For example, a variable $x : \Pi(y:W_A).W_B$ is normal, but can be η -expanded to $\lambda(y:W_A).x y$. In order to properly η -expand a term, we need additional type information, which our semantic model will include.

3.2.1 Untyped Domain Model. Following [Abel \[2013\]](#) we give the definition of the untyped domain model. It is the target to which well-typed terms are evaluated to.

De Bruijn levels	$\mathbb{N}$	$\ni$	z	
Neutral domain values	D^{Ne}	$\ni$	v	$::= l_z \mid \text{rec}(A, b, M_{\text{succ}}, \rho, v) \mid v \ w$
Domain values	D	$\ni$	a, b, m, n	$::= \uparrow^a(v) \mid N \mid U@i \mid \Pi(a, B, \rho)$ $\mid \text{ze} \mid \text{su}(a) \mid \Lambda(M, \rho)$
Normal domain values	D^{Nf}	$\ni$	w	$::= \downarrow^a(m)$
Evaluation environments	Env	$\ni$	ρ	

In contrast to the syntax's de Bruijn indices, domain variables are represented with de Bruijn levels. This means that we count from the left instead of the right, resulting in every variable having a unique de Bruijn level throughout a program, thus removing the need to shift indices. De Bruijn indices and levels have a correspondence. In general, given a context $\Gamma, x:T, \Delta$, the de Bruijn index d of x is $|\Delta|$, the length of Δ , whereas the de Bruijn level z of x is $|\Gamma|$. Therefore, we have $d + z + 1 = |\Gamma, x : T, \Delta|$. This correspondence will be used later in the readback function to read de Bruijn indices back from de Bruijn levels. Otherwise, the domain is defined similarly to the (syntactic) neutral and normal forms.

The model is split into three definitions. Neutral domain values are just domain variables (l_z) or elimination forms. They correspond closely to syntactic neutral forms. Domain values can be reflected neutral values annotated with a domain type, $\uparrow^a(v)$, or introduction forms. They correspond closely to the syntactic normal forms, with one major difference: to mechanize dependent functions in the domain values, we employ defunctionalization [[Reynolds 1998](#)] to capture the ambient evaluation environment ρ and an open type B or an open term M . This is also standard [[Abel 2013](#); [Hu et al. 2023](#)]. Finally, normal domain values are reified domain values annotated with a domain type, $\downarrow^a(m)$. This represents the η -long form of m , although $\downarrow^a(m)$ is not itself η -long since reification is simply a marker. Rather, we will later use a to guide the η -expansions that need to be performed during readback.

Next, an environment ρ is just a function from de Bruijn indices to D . We need a number of tools related to evaluation environments. We omit their concrete definitions for space consideration. First, the empty environment $\text{emp} :: \text{Env}$ interprets the empty context as a constant function. We use ze as the default value of emp , but the concrete value does not actually matter. Later, we will prove the soundness property, which requires that all variables are bound and implies that we will never access the default value at all. Then, we also need two useful operations on environments: The extension function $\text{ext}(\rho, m)$ shifts everything in ρ up by 1 and produces m on 0; its inverse drop function $\text{drop}(\rho)$ shifts everything in ρ down by 1, forgetting about the previous 0^{th} value.

3.3 Normalization by Evaluation for MLTT

Next, we define the normalization procedure. It consists of two main steps. First, we define evaluation functions that map syntactic objects to semantic objects, eliminating any β -redexes in the process. Second, we define readback functions that map semantic objects back to syntactic objects, performing all the necessary η -expansions in the process. Then, we can compose evaluation and readback to obtain a normalization procedure for computing $\beta\eta$ -normal forms.

3.3.1 Evaluation. As the first step of the NbE algorithm, we evaluate a well-typed term into a domain value. Evaluation is a partial function (denoted by $\rightarrow$). Since RocQ requires all functions to be total, our implementation models evaluation as functional relations.

$$\begin{aligned}
& \llbracket _ \rrbracket (_) :: \text{Exp} \rightarrow \text{Env} \rightarrow \text{D} \\
& \llbracket x_d \rrbracket (\rho) := \rho(d) \\
& \llbracket \Pi(x : A).B \rrbracket (\rho) := \Pi(\llbracket A \rrbracket (\rho), B, \rho) \\
& \llbracket \lambda(x : A).M \rrbracket (\rho) := \Lambda(M, \rho) \\
& \llbracket M N \rrbracket (\rho) := \llbracket M \rrbracket (\rho) \cdot \llbracket N \rrbracket (\rho) \\
& \llbracket t[\sigma] \rrbracket (\rho) := \llbracket t \rrbracket (\llbracket \sigma \rrbracket (\rho)) \\
\\
& _ \cdot _ :: \text{D} \rightarrow \text{D} \rightarrow \text{D} \\
& (\Lambda(M, \rho)) \cdot a := \llbracket M \rrbracket (\text{ext}(\rho, a)) \\
& (\uparrow^{\Pi(a, B, \rho)}(v)) \cdot m := \uparrow^{\llbracket B \rrbracket (\text{ext}(\rho, m))}(v \downarrow^a(m))
\end{aligned}$$

Fig. 3. Evaluation functions

We define four partial functions for the whole evaluation procedure.

Evaluation of expressions	$\llbracket _ \rrbracket (_) :: \text{Exp} \rightarrow \text{Env} \rightarrow \text{D}$
Evaluation of substitutions	$\llbracket _ \rrbracket_s (_) :: \text{Subst} \rightarrow \text{Env} \rightarrow \text{Env}$
Domain application	$_ \cdot _ :: \text{D} \rightarrow \text{D} \rightarrow \text{D}$
Domain recursion for Nat	$\text{rec} \cdot (_ , _ , _ , _) :: \text{Exp} \rightarrow \text{D} \rightarrow \text{Exp} \rightarrow \text{Env} \rightarrow \text{D} \rightarrow \text{D}$

The first function is the main one: $\llbracket M \rrbracket (\rho)$ performs evaluation of a syntactic expression M in the evaluation environment ρ . The second function performs evaluation of substitutions into environments. This is required since we use explicit substitutions. The last two functions perform computations inside the domain: $m \cdot n$ is function application and $\text{rec} \cdot (A, m_{\text{zero}}, M_{\text{succ}}, \rho, m)$ is recursion over natural numbers. We show the definition of the main evaluation function together with the Domain application in Fig. 3. The complete definition is available in Appendix A (and [its mechanized documentation](#)).

Evaluation is a straightforward process: Every object is evaluated by applying congruences as much as possible, then, on elimination forms, we apply the relevant domain computation. One subtlety occurs in the treatment of open objects, such as functions or the successor branch of the recursor. Specifically, these correspond to premises that are in a larger context than their conclusion, which means that our current evaluation environment ρ does not yet provide an interpretation of the additional variables. Thus, we cannot evaluate them until a later stage when all the necessary information is available. In the meantime, we exploit defunctionalization and simply keep track of the syntactic forms and of the current evaluation environment.

3.3.2 Readback. In the second step of the NbE algorithm, we read the domain values back to a syntactic normal form while performing any necessary η -expansion. The readback process consists of three mutually defined partial functions (see Fig. 4), which are again defined as functional relations in Rocq. First, R^{Nf} reads back a normal domain object to a $\beta\eta$ -normal syntactic object. R^{Nf} is the main function and it is responsible for performing the η -expansion. Second, R^{Ty} reads back a type in $\beta\eta$ -normal form. R^{Ty} is in particular used to read back domain objects of the form $\downarrow^{U@i}(a)$. Third, R^{Ne} reads back neutral domain objects to neutral syntactic objects. R^{Ne} is responsible for converting between de Bruijn levels and indices and it is largely defined by congruence rules.

The readback functions take in addition a natural number z as a parameter, which measures the length of current typing context. This number is used in the variable case to convert from de Bruijn levels into de Bruijn indices:

$$R^{\text{Ne}}_{z'}(l_z) := x_{\max(z'-z-1, 0)}$$

The formula has been explained at the beginning of Sec. 3.2.1.

$$\begin{aligned}
& \mathbf{R}_{-}^{\text{Nf}}(_) :: \mathbb{N} \rightarrow \mathbf{D}^{\text{Nf}} \rightarrow \text{Nf} \\
& \mathbf{R}_z^{\text{Nf}}(\downarrow^a (\uparrow^b (v))) := \mathbf{R}_z^{\text{Ne}}(v) \quad (\text{where } a = \uparrow^{a'}(c) \text{ or } \mathbf{N}) \\
& \mathbf{R}_z^{\text{Nf}}(\downarrow^{\text{U@}i} (a)) := \mathbf{R}_z^{\text{Ty}}(a) \\
& \mathbf{R}_z^{\text{Nf}}(\downarrow^{\Pi(a,B,\rho)} (m)) := \lambda(x : \mathbf{R}_z^{\text{Ty}}(a)). \mathbf{R}_{1+z}^{\text{Nf}}(\downarrow^{\llbracket B \rrbracket(\text{ext}(\rho, \uparrow^a(l_z)))} (m \cdot \uparrow^a(l_z))) \\
& \mathbf{R}_{-}^{\text{Ty}}(_) :: \mathbb{N} \rightarrow \mathbf{D} \rightarrow \text{Nf} \\
& \mathbf{R}_z^{\text{Ty}}(\Pi(a, B, \rho)) := \Pi(x : \mathbf{R}_z^{\text{Ty}}(a)). \mathbf{R}_{1+z}^{\text{Ty}}(\llbracket B \rrbracket(\text{ext}(\rho, \uparrow^a(l_z)))) \\
& \mathbf{R}_z^{\text{Ty}}(\uparrow^a(v)) := \mathbf{R}_z^{\text{Ne}}(v) \\
& \mathbf{R}_{-}^{\text{Ne}}(_) :: \mathbb{N} \rightarrow \mathbf{D}^{\text{Ne}} \rightarrow \text{Ne} \\
& \mathbf{R}_{z'}^{\text{Ne}}(l_z) := x_{\max(z'-z-1, 0)} \\
& \mathbf{R}_z^{\text{Ne}}(v \ w) := \mathbf{R}_z^{\text{Ne}}(v) \ \mathbf{R}_z^{\text{Nf}}(w)
\end{aligned}$$

Fig. 4. Definition of readbacks

3.3.3 Normalization. With the evaluation and the readback functions, we are able to provide the NbE algorithm. First, we define a helper function to compute the initial environment:

$$\begin{aligned}
& \uparrow :: \text{Ctx} \rightarrow \text{Env} \\
& \uparrow \cdot := \text{emp} \\
& \uparrow^{\Gamma, x:A} := \text{ext}(\uparrow^{\Gamma}, \uparrow^{\llbracket A \rrbracket(\uparrow^{\Gamma})} (l_{|\Gamma|}))
\end{aligned}$$

Then the NbE algorithm is defined by first evaluating the term and the type, and then reading back the domain value of the term annotated by the domain value of the type.

Definition 3.1. For $\Gamma \vdash M : A$, the NbE algorithm for terms is

$$\mathbf{nbe}_{\Gamma}^A(M) := \mathbf{R}_{|\Gamma|}^{\text{Nf}}(\downarrow^{\llbracket A \rrbracket(\uparrow^{\Gamma})} (\llbracket M \rrbracket(\uparrow^{\Gamma})))$$

We need another variant of the NbE algorithm, which normalizes a type. The steps are similar, but we use $\mathbf{R}^{\text{Ty}}$ for readback, because we know we are handling a type.

Definition 3.2. For $\Gamma \vdash A : \text{Type@}i$, the NbE algorithm for types is

$$\mathbf{nbe}_{\Gamma}(A) := \mathbf{R}_{|\Gamma|}^{\text{Ty}}(\llbracket A \rrbracket(\uparrow^{\Gamma}))$$

4 Completeness and Soundness of NbE

In this section, we establish that our NbE is a complete and sound algorithm. In this setting, completeness means that NbE gives the same normal form for two equivalent input expressions. On the other hand, soundness means that the normal form out of NbE is equivalent to the input expression. Both properties together imply several important colloraries on typing, subtyping, and equivalence: decidability of equivalence, decidability of subtyping, and completeness and soundness of algorithmic subtyping and type-checking.

However, as the NbE algorithm crosses through a domain value, we need to define a relation model between two domain types/values and another model between a domain value and an expression to show the completeness and soundness.

4.1 Partial Equivalence Relation (PER) Model

We first need to provide the partial equivalence relation (PER) model, which relates two domain types and domain values. In previous work [Abel 2013; Hu et al. 2023], the PER model is defined using induction-recursion [Dybjer and Setzer 2003; Martin-Löf 1984]. However, Rocq does not

support induction-recursion. As an alternative, we use the impredicative Prop universe in Rocq. In this universe, we can define a proposition by quantifying over all propositions in Prop, including the one being defined. This feature augments the logical power of Rocq enough to express the PER model of a dependent type theory. In fact, [Wieczorek and Biernacki \[2018\]](#) use Prop to capture induction-recursion to mechanize NbE of MLTT by adapting the Bove-Capretta method [\[Bove 2009; Bove and Capretta 2005\]](#). In this approach, their PER model is unrolled into two separate relations, one for domain types and one for domain values. However, this is only possible for a bounded universe hierarchy, and indeed their mechanization is restricted to one universe. This limitation of applying the Bove-Capretta method to universe hierarchy has been noticed by [Abel et al. \[2017\]](#), but their solution is based on sized types, which are not available in Rocq either.

The key intuition of our solution is that we do not need computability of the recursive definition. Previous methods [\[Bove 2009; Bove and Capretta 2005; Larchey-Wendling and Monin 2018\]](#) focus on a faithful encoding of induction-recursion that keeps computability of the recursive function. However, in our case, we use the PER model only to prove properties of NbE and not to define the NbE algorithm. Thus, we can encode the inductive-recursive PER model into a (weak partial) functional relation, which does not directly provide computability. This functional relation is defined inductively and relates two domain types in the universe of a given level. Then, as an output, it “returns” a binary relation between domain values of those related domain types.

To define the PER model, we need a few more auxiliary definitions. First, we define three PERs based on readback, $Nf \subseteq D^{Nf} \times D^{Nf}$, $Ne \subseteq D^{Ne} \times D^{Ne}$, and $Ty \subseteq D \times D$:

$$\frac{\forall z. R_z^{Nf}(w) = R_z^{Nf}(w')}{w \approx w' \in Nf} \quad \frac{\forall z. R_z^{Ne}(v) = R_z^{Ne}(v')}{v \approx v' \in Ne} \quad \frac{\forall z. R_z^{Ty}(a) = R_z^{Ty}(a')}{a \approx a' \in Ty}$$

Nf relates two normal domain values if and only if their readbacks are equal in any given context. Similarly, Ne relates two neutral domain values based on readbacks, and Ty relates two domain types based on readbacks. We use these PERs as base cases for the PER model. Furthermore, we describe the realizability of the PER model using these PERs, which leads to the completeness theorem of NbE. Therefore, they also play the key roles for the completeness theorem as well.

Now, we first define the PER model $Neu \subseteq D \times D$ for domain values of neutral domain types.

$$\frac{w \approx w' \in Ne}{\uparrow^a(w) \approx \uparrow^{a'}(w') \in Neu}$$

This relation Neu simply ignores the type annotations and relates two neutral domain values using Ne . We can safely ignore these type annotations as we will use Neu only for well-typed terms and check the relation between types first with our PER model for domain types. Next, we define the PER model $Nat \subseteq D \times D$ for domain values of the natural number type.

$$\frac{}{ze \approx ze \in Nat} \quad \frac{m \approx n \in Nat}{su(m) \approx su(n) \in Nat} \quad \frac{v \approx v' \in Ne}{\uparrow^a(v) \approx \uparrow^{a'}(v') \in Nat}$$

This relates ze with itself, $su(-)$ of related natural numbers, and two neutral domain values. Again, the case of neutral domain values ignores the type annotations.

Now, we can define the PER model $\mathcal{U}_i$ for domain types in the i -th universe. We write

$$a \approx a' \in \mathcal{U}_i \searrow R$$

to relate domain types a and a' , and a binary relation R relating two domain values of types a and a' . The relation R is regarded as a “return” predicate, so $\mathcal{U}_i$ is a functional relation. We write $a \approx a' \in \mathcal{U}_i$ if we are only concerned about the relation between a and a' . The definition is given in Fig. 5. There are 4 cases: one for types in neutral forms and one for each type constructor

$$\begin{array}{c}
\frac{b \approx b' \in Ne \quad R \Leftrightarrow Neu}{\uparrow^a(b) \approx \uparrow^{a'}(b') \in \mathcal{U}_i \searrow R} \quad \frac{j < i \quad R \Leftrightarrow \mathcal{U}_j}{U@j \approx U@j \in \mathcal{U}_i \searrow R} \quad \frac{R \Leftrightarrow Nat}{N \approx N \in \mathcal{U}_i \searrow R} \\
\\
\frac{
\begin{array}{c}
a \approx a' \in \mathcal{U}_i \searrow R_I \\
R_I \text{ is a PER} \quad \forall \mathcal{D} :: m \approx m' \in R_I . \llbracket B \rrbracket(\text{ext}(\rho, m)) \approx \llbracket B' \rrbracket(\text{ext}(\rho', m')) \in \mathcal{U}_i \searrow R_O(\mathcal{D}) \\
\forall n, n' . n \approx n' \in R \leftrightarrow (\forall \mathcal{D} :: m \approx m' \in R_I . n \cdot m \approx n' \cdot m' \in R_O(\mathcal{D}))
\end{array}
}{\Pi(a, B, \rho) \approx \Pi(a', B', \rho') \in \mathcal{U}_i \searrow R}
\end{array}$$

Fig. 5. PER Model for Domain Types and Values

$$\begin{array}{c}
\frac{b \approx b' \in Ne}{\uparrow^a(b) <_i \uparrow^{a'}(b')} \quad \frac{j \leq k < i}{U@j <_i U@k} \quad \frac{}{N <_i N} \\
\\
\frac{
\begin{array}{c}
a \approx a' \in \mathcal{U}_i \searrow R_I \quad \forall \mathcal{D} :: m \approx m' \in R_I . \llbracket B \rrbracket(\text{ext}(\rho, m)) <_i \llbracket B' \rrbracket(\text{ext}(\rho', m')) \\
\Pi(a, B, \rho) \approx \Pi(a, B, \rho) \in \mathcal{U}_i \quad \Pi(a', B', \rho') \approx \Pi(a', B', \rho') \in \mathcal{U}_i
\end{array}
}{\Pi(a, B, \rho) <_i \Pi(a', B', \rho')}
\end{array}$$

Fig. 6. Semantic Subtyping

($U@j$, N , and Π). In the Π case, we have one **highlighted** premise. This premise enhances the induction principle for $\mathcal{U}$ so that we can prove $\mathcal{U}$ and its return relation R are actually PERs without postulating any axioms. After we prove that R is actually PER, we remove this **highlighted** premise as it is logically redundant.

It is worth noting that we define $\mathcal{U}$ as a *weak* partial functional relation, in the sense that the return relation R is only unique up to relational equivalence. Here, the relational equivalence $P \Leftrightarrow Q$ between two n -ary relations P and Q is defined as

$$\forall x_1, x_2, \dots, x_n . P \ x_1 \ x_2 \ \dots \ x_n \leftrightarrow Q \ x_1 \ x_2 \ \dots \ x_n$$

We employ this weakness to avoid excessive uses of extensionality axioms, especially propositional and functional extensionality. In fact, this approach allows us to remove propositional extensionality entirely, and only keep one use of functional extensionality for convenience.

Given the definition, it becomes clear why $\mathcal{U}$ is encoded in the impredicative universe Prop . In the universe case, $\mathcal{U}_i$ returns R , a relation equivalent to $\mathcal{U}_j$ for $j < i$. In a predicative setting, this would have required $\mathcal{U}_i$ to live in a higher universe than $\mathcal{U}_j$, but this is not possible because they are the same inductive definition. In Prop , this problem no longer exists thanks to impredicativity.

In addition to the PER model, since we are working with universe subtyping in the syntactic judgements, we also need to introduce subtyping to the semantics as well. The semantic subtyping relation $a <_i a'$ between two domain types a and a' at universe level i is defined in Fig. 6. The four rules in semantic subtyping reflects the definition of syntactic subtyping. For example, in the Π case, the input types a and a' are semantically equivalent by $\mathcal{U}$, because we employ covariant subtyping in our syntactic rules.

Finally, the PER model and semantic subtyping are extended to PER and semantic subtyping between contexts and environments in order to define the semantic typing judgement for the completeness theorem. We list the cons rule of the PER $\Gamma \approx \Gamma' \searrow R$ between contexts Γ and Γ' and the cons rule of semantic subtyping between contexts $\Gamma <: \Gamma'$ in Fig. 7. We omit the rules for the empty context since they simply relate the empty context to itself. Note that the PER between contexts also gives the PER R for environments as a “returned” value, similarly to the definition of $\mathcal{U}$. In case of the empty context, this returned PER is the total relation, i.e. any two environments are related under the PER. Another subtle but important point in the definition is that universe

$$\begin{array}{c}
\frac{\Gamma \approx \Gamma' \searrow R_T \quad \boxed{R_T \text{ is a PER}} \quad \forall \mathcal{D} :: \rho \approx \rho' \in R_T . \llbracket A \rrbracket(\rho) \approx \llbracket A' \rrbracket(\rho') \in \mathcal{U}_i \searrow R_H(\mathcal{D})}{\Gamma, x:A \approx \Gamma', x:A' \searrow R} \\
\frac{\forall \mathcal{D} :: \rho \approx \rho' \in R_T . \llbracket A \rrbracket(\rho) <_i \llbracket A' \rrbracket(\rho') \quad \Gamma <: \Gamma' \quad \Gamma \approx \Gamma' \searrow R_T \quad \Gamma, x:A \approx \Gamma, x:A \searrow R_1 \quad \Gamma', x:A' \approx \Gamma', x:A' \searrow R_2}{\Gamma, x:A <: \Gamma', x:A'}
\end{array}$$

Fig. 7. The Cons Rule of the PER Model and Semantic Subtyping for Contexts

level i , in which we relate the heads $\llbracket A \rrbracket(\rho)$ and $\llbracket A' \rrbracket(\rho')$, is hidden in the conclusion. This forces us to prove a slightly more general property about $\mathcal{U}$ in order to prove that the PER model for contexts is indeed transitive in Sec. 4.2.2.

In the following section, we prove the properties of the PER model for domain types and values leading to the completeness theorem.

4.2 Properties of the PER Model

4.2.1 Irrelevance. To show the interesting properties of the PER models, we first need to specify what determines R in the definition of $\mathcal{U}$. Our intuition is that we do not need a full derivation of $a \approx b \in \mathcal{U}_i \searrow R$ to determine R , but having a domain type a that satisfies $a \approx b \in \mathcal{U}_i$ is enough. The following lemma formalizes this intuition:

LEMMA 4.1 (**U RIGHT IRRELEVANCE**). *If $a \approx b \in \mathcal{U}_i \searrow R$ and $a \approx b' \in \mathcal{U}_i' \searrow R'$, then $R \Leftrightarrow R'$.*

Thus, as long as we have the domain type a , the relation between domain values of that type is unique up to relational equivalence.

4.2.2 Structural Property of the PER Model. Now, we show the structural properties of the PER model and semantic subtyping: the PER model is indeed a PER and semantic subtyping is transitive. We start by showing that the PER model is symmetric, both for domain types and values.

LEMMA 4.2 (**SYMMETRY OF U**). *If $a \approx a' \in \mathcal{U}_i \searrow R$, then*

- $a' \approx a \in \mathcal{U}_i \searrow R$, and,
- $m \approx m' \in R$ implies $m' \approx m \in R$.

With this, we can show that it is also transitive.

LEMMA 4.3 (**TRANSITIVITY OF U**). *If $a_1 \approx a_2 \in \mathcal{U}_i \searrow R$, then*

- $a_2 \approx a_3 \in \mathcal{U}_j \searrow R$ implies $a_1 \approx a_3 \in \mathcal{U}_i \searrow R$, and,
- $m_1 \approx m_2 \in R$ and $m_2 \approx m_3 \in R$ implies $m_1 \approx m_3 \in R$

Note that this theorem is stronger than the naïve transitivity because we can choose a different universe level j for $a_2 \approx a_3 \in \mathcal{U}_j \searrow R$. This flexibility is the key for the transitivity of the PER for contexts, as three transitively related contexts might use different universe levels to relate their heads. The transitivity of the semantic subtyping for domain types, the symmetry and the transitivity of the PER model for contexts and environments, and the transitivity of the semantic subtyping for contexts follow from the above two in this order.

4.2.3 Realizability of the PER Model. Another important property of $\mathcal{U}$ is realizability, which is essential to establish both completeness and soundness. We define the realizability using Nf , Ne , and Ty following Abel [2013]:

THEOREM 4.4 (REALIZABILITY OF PER MODEL). *Given $a \approx a' \in \mathcal{U}_i \searrow R$,*

- *For domain values:* $m \approx m' \in R$ implies $\downarrow^a(m) \approx \downarrow^{a'}(m') \in Nf$;
- *For neutral domain values:* $w \approx w' \in Ne$ implies $\uparrow^a(w) \approx \uparrow^{a'}(w') \in R$;
- *For domain types:* $a \approx a' \in Ty$.

PROOF. By induction on $a \approx a' \in \mathcal{U}_i \searrow R$. We need to prove a lemma that $m \approx m' \in Nat$ implies $\downarrow^N m \approx \downarrow^N m' \in Nf$ along the way, which is shown by induction on $m \approx m' \in Nat$. $\square$

4.2.4 Other Properties of the PER Model. To show the fundamental theorem for completeness (that syntactic typing rules are admissible for semantic equivalence judgement for completeness), we need some lemmas about $\mathcal{U}$. Particularly, we need that $\mathcal{U}$ is cumulative in its universe level.

LEMMA 4.5 (CUMULATIVITY OF $\mathcal{U}$). *If $a \approx a' \in \mathcal{U}_i \searrow R$, then $a \approx a' \in \mathcal{U}_j \searrow R$ for $j \geq i$.*

Another key property is that the PER respects subtyping, which is crucial for our subtyping rule.

LEMMA 4.6 (SUBTYPING OF $\mathcal{U}$). *If $a <_i a'$, $a \approx b \in \mathcal{U}_i \searrow R$, and $a' \approx b' \in \mathcal{U}_i \searrow R'$ then $m \approx m' \in R$ implies $m \approx m' \in R'$.*

4.3 Completeness

We first define the semantic equivalence judgement $\Gamma \models M \approx M' : A$ for completeness:

$$\frac{\Gamma \approx \Gamma' \searrow R_E \quad \forall \mathcal{D} :: \rho \approx \rho' \in R_E \cdot \llbracket A \rrbracket(\rho) \approx \llbracket A \rrbracket(\rho') \in \mathcal{U}_i \searrow R \wedge \llbracket M \rrbracket(\rho) \approx \llbracket M' \rrbracket(\rho') \in R}{\Gamma \models M \approx M' : A}$$

From this, we derive the semantic typing judgement $\Gamma \models M : A$ by defining it as an alias of $\Gamma \models M \approx M : A$. Then, we use the lemmas for $\mathcal{U}$ to show the following fundamental theorem of the semantic judgements for completeness:

THEOREM 4.7 (FUNDAMENTAL THEOREM). *If $\Gamma \vdash M \approx M' : A$ then $\Gamma \models M \approx M' : A$.*

PROOF. By induction on the given equivalence derivation. This should be mutually proven with fundamental theorems for other judgements such as one for $\Gamma \vdash M : A$. In our mechanization, we handle each case as a separate lemma as they are fairly long. $\square$

Now, we can derive the completeness of NbE:

THEOREM 4.8 (COMPLETENESS OF NbE). *If $\Gamma \vdash M \approx M' : A$ then $nbe_\Gamma^A(M) = nbe_\Gamma^A(M')$.*

PROOF. We first show a lemma that $\Gamma \approx \Gamma' \searrow R_E$ implies $\uparrow^\Gamma \approx \uparrow^{\Gamma'} \in R_E$ by induction on the given derivation. Then, we prove the main goal by applying the fundamental theorem together with this lemma. $\square$

4.4 Kripke Gluing Model

Now, we move our focus to the soundness of NbE. For soundness, we need to connect a term and its normalization result that was read back from the semantic world. Thus, the model for soundness should “glue” the syntactic types/terms with domain types/values [Coquand and Dybjer 1997]. Moreover, as we need to deal with open terms, our model should be stable under weakenings. Following Abel [2013], we call this condition “Kripke”. In short, we will define a Kripke gluing model for McTT in this section.

First, we need to formalize the definition of weakenings. The following judgement $\Gamma \vdash_w \sigma : \Delta$ means that σ is a weakening from Δ to Γ .

$$\frac{\Gamma \vdash \sigma \approx \text{id} : \Delta}{\Gamma \vdash_w \sigma : \Delta} \quad \frac{\Gamma \vdash_w \tau : \Delta', x:A \quad \vdash \Delta' \subseteq \Delta \quad \Gamma \vdash \sigma \approx \text{wk} \circ \tau : \Delta}{\Gamma \vdash_w \sigma : \Delta}$$

$$\begin{array}{c}
\frac{P \Leftrightarrow \text{Neu}_i^a \quad El \Leftrightarrow \text{Neu}_i^a}{\uparrow^b a \in \mathcal{U}_i^* \searrow P \searrow El} \quad \frac{P \Leftrightarrow \text{Nat}_i \quad El \Leftrightarrow \text{Nat}_i}{N \in \mathcal{U}_i^* \searrow P \searrow El} \\
j < i \quad \forall \Gamma, A. \Gamma \vdash A \textcircled{R} P \leftrightarrow (\Gamma \vdash A \approx \text{Type}@j : \text{Type}@i) \\
\frac{\forall \Gamma, M, A, m. \Gamma \vdash M : A \textcircled{R} m \in El \leftrightarrow (\Gamma \vdash M : A \wedge \Gamma \vdash A \textcircled{R} P \wedge m \in \mathcal{U}_j^* \searrow P' \searrow El' \wedge \Gamma \vdash M \textcircled{R} P')}{U@j \in \mathcal{U}_i^* \searrow P \searrow El} \\
\frac{a \in \mathcal{U}_i^* \searrow P_I \searrow El_I \quad a \approx a \in \mathcal{U}_i \searrow R_I \quad \forall \mathcal{D} :: n \approx n \in R_I. \llbracket B \rrbracket(\text{ext}(\rho, n)) \in \mathcal{U}_i^* \searrow P_O(\mathcal{D}) \searrow El_O(\mathcal{D})}{\Pi(a, B, \rho) \approx \Pi(a, B, \rho) \in \mathcal{U}_i \searrow R \quad P \Leftrightarrow \Pi_i(R_I, P_I, El_I, P_O) \quad El \Leftrightarrow \Pi_i(R, R_I, P_I, El_I, El_O)} \\
\frac{}{\Pi(a, B, \rho) \in \mathcal{U}_i^* \searrow P \searrow El}
\end{array}$$

Fig. 8. Gluing Model for Domain Types and Values

Note that the rule for wk also handles context subtyping. This allows us to apply subtyping between contexts to weakening judgement without a separate rule for subtyping.

Now, we define the Kripke gluing model. Usually, one defines the model using a recursion over a derivation of the PER model for domain types. However, our $\mathcal{U}$ is defined in Prop, and we cannot match on it to generate a new predicate.³ Thus, we need to define an inductive predicate $\mathcal{U}^*$ reusing cases of $\mathcal{U}$. There are two main differences between $\mathcal{U}^*$ and $\mathcal{U}$: 1) $\mathcal{U}^*$ is a unary relation on domain types, while $\mathcal{U}$ is binary; 2) $\mathcal{U}^*$ outputs two predicates, while $\mathcal{U}$ only outputs one. More precisely, the judgement $a \in \mathcal{U}_i^* \searrow P \searrow El$ returns a unary predicate P for syntactic types related to a and a binary predicate El for syntactic terms related to a domain value of type a . To give a clear definition of $\mathcal{U}^*$, we first need to specify concrete predicates P and El . We will use notations $\Gamma \vdash A \textcircled{R} P$ for gluing of a syntactic type A , and $\Gamma \vdash M : A \textcircled{R} m \in El$ for gluing of the syntactic term M and domain value m .

For the neutral case, we define Neu_i^a that relates a syntactic type A under a context Γ with a given neutral domain type a and universe level i ($\Gamma \vdash A \textcircled{R} \text{Neu}_i^a$). For brevity, Neu_i^a is also overloaded for the relation that relates a syntactic term M of type A with a domain value m ($\Gamma \vdash M : A \textcircled{R} m \in \text{Neu}_i^a$).

$$\begin{array}{c}
\frac{\Gamma \vdash A : \text{Type}@i \quad \forall \Delta, \sigma. \Delta \vdash_w \sigma : \Gamma \rightarrow \Delta \vdash A[\sigma] \approx R_{|\Delta|}^{\text{Ne}}(a) : \text{Type}@i}{\Gamma \vdash A \textcircled{R} \text{Neu}_i^a} \\
\frac{\Gamma \vdash A \textcircled{R} \text{Neu}_i^a \quad \Gamma \vdash M : A \quad v \approx v \in Ne \quad \forall \Delta, \sigma. \Delta \vdash_w \sigma : \Gamma \rightarrow \Delta \vdash M[\sigma] \approx R_{|\Delta|}^{\text{Ne}}(v) : \text{Type}@i}{\Gamma \vdash M : A \textcircled{R} \uparrow^{a'}(v) \in \text{Neu}_i^a}
\end{array}$$

Note that both $\Gamma \vdash A \textcircled{R} \text{Neu}_i^a$ and $\Gamma \vdash M : A \textcircled{R} m \in \text{Neu}_i^a$ require domain and syntax to be related even after an arbitrary weakening from Γ . We repeat this requirement for the N case whenever neutral terms appear. We use the name Nat_i for the relation that checks syntactic type A is a natural number type ($\Gamma \vdash A \textcircled{R} \text{Nat}_i$) and for the relation between syntactic natural number term M with domain natural number m ($\Gamma \vdash M : A \textcircled{R} m \in \text{Nat}_i$). This relation needs to be inductive, and we

³It is possible to match on Prop to generate a proof of a predicate of Prop. However, Prop itself lives in Type, and one cannot generate a member of a type by pattern matching on Prop.

$$\begin{array}{c}
\frac{\forall \Delta, \sigma, \rho . \Delta \vdash \sigma \textcircled{R} \rho \in Sb \leftrightarrow \Delta \vdash \sigma : \cdot}{\cdot \searrow Sb} \\
\\
\frac{\Gamma \searrow Sb_T \quad \Gamma \vdash A : \text{Type}@i \quad Sb \Leftrightarrow \text{Cons}_i(\Gamma, A, Sb_T) \quad \forall \Delta, \sigma, \rho, P, El . \Delta \vdash \sigma \textcircled{R} \rho \in Sb_T \rightarrow \llbracket A \rrbracket(\rho) \in \mathcal{U}_i^* \searrow P \searrow El \wedge \Delta \vdash A[\sigma] \textcircled{R} P}{\Gamma, x:A \searrow Sb}
\end{array}$$

Fig. 9. Gluing Model for Contexts

extract the inductive part into $\Gamma \vdash M \textcircled{R} m \in \text{Nat}$.

$$\begin{array}{c}
\frac{\Gamma \vdash A \approx \text{Nat} : \text{Type}@i}{\Gamma \vdash A \textcircled{R} \text{Nat}_i} \quad \frac{\Gamma \vdash A \textcircled{R} \text{Nat}_i \quad \Gamma \vdash M \textcircled{R} m \in \text{Nat}}{\Gamma \vdash M : A \textcircled{R} m \in \text{Nat}_i} \quad \frac{\Gamma \vdash M \approx \text{zero} : \text{Nat}}{\Gamma \vdash M \textcircled{R} \text{ze} \in \text{Nat}} \\
\\
\frac{\Gamma \vdash M \approx \text{succ } M' : \text{Nat} \quad \Gamma \vdash M' \textcircled{R} m' \in \text{Nat}}{\Gamma \vdash M \textcircled{R} \text{su}(m') \in \text{Nat}} \quad \frac{v \approx v \in \text{Ne} \quad \forall \Delta, \sigma . \Delta \vdash_w \sigma : \Gamma \rightarrow \Delta \vdash M[\sigma] \approx R_{|\Delta|}^{\text{Ne}}(v) : \text{Type}@i}{\Gamma \vdash M \textcircled{R} \uparrow^a(w) \in \text{Nat}}
\end{array}$$

Likewise, we define $\Pi_i(R_I, P_I, El_I, P_O)$ for Π -types and $\Pi_i(R, R_I, P_I, El_I, El_O)$ for terms of Π -types, where the subscripts I and O indicate a relation on inputs and outputs, respectively, and R is the PER associated to the whole function space:

$$\begin{array}{c}
\frac{\Gamma \vdash A \approx \Pi(x:B).C : \text{Type}@i \quad \forall \Delta, \sigma . \Delta \vdash_w \sigma : \Gamma \rightarrow \Delta \vdash B[\sigma] \textcircled{R} P_I \quad \forall \Delta, \sigma, N, n . \Delta \vdash_w \sigma : \Gamma \wedge \Delta \vdash N : B[\sigma] \textcircled{R} n \in El_I \rightarrow \forall \mathcal{D} :: n \approx n \in R_I . \Delta \vdash C[\sigma, N] \textcircled{R} P_O(\mathcal{D})}{\Gamma \vdash A \textcircled{R} \Pi_i(R_I, P_I, El_I, P_O)} \\
\\
\frac{\Gamma \vdash M : A \quad m \approx m \in R \quad \Gamma \vdash A \approx \Pi(x:B).C : \text{Type}@i \quad \forall \Delta, \sigma . \Delta \vdash_w \sigma : \Gamma \rightarrow \Delta \vdash B[\sigma] \textcircled{R} P_I \quad \forall \Delta, \sigma, N, n . \Delta \vdash_w \sigma : \Gamma \wedge \Delta \vdash N : B[\sigma] \textcircled{R} n \in El_I \rightarrow \forall \mathcal{D} :: m \approx m \in R_I . \Delta \vdash M[\sigma] N : C[\sigma, N] \textcircled{R} m \cdot n \in El_O(\mathcal{D})}{\Gamma \vdash M : A \textcircled{R} m \in \Pi_i(R, R_I, P_I, El_I, El_O)}
\end{array}$$

Using these predicates, we give the definition of $\mathcal{U}^*$ in Fig. 8.

We also need to extend the gluing model to one for contexts for soundness. In Fig. 9, we define $\Gamma \searrow Sb$. Here, Sb is a predicate relating a syntactic substitution σ from Δ to Γ with an environment ρ . To represent this, we use the notation $\Delta \vdash \sigma \textcircled{R} \rho \in Sb$. This definition uses the following predicate $\text{Cons}_i(\Gamma, A, Sb_T)$:

$$\frac{\Delta \vdash \sigma : \Gamma, x:A \quad \llbracket A \rrbracket(\rho) \in \mathcal{U}_i^* \searrow P \searrow El \quad \Delta \vdash x[\sigma] : A[\text{wk} \circ \sigma] \textcircled{R} \rho(0) \in El \quad \Delta \vdash \text{wk} \circ \sigma \textcircled{R} \text{drop}(\rho) \in Sb_T}{\Gamma \vdash \sigma \textcircled{R} \rho \in \text{Cons}_i(\Gamma, x, A, Sb_T)}$$

This predicate relates the substitution entry for latest variable in the context ($x[\sigma]$) with the top-most value in the environment ($\rho(0)$), and relate weakened substitution ($\text{wk} \circ \sigma$) with the shifted environment ($\text{drop}(\rho)$).

In the following section, we prove the properties of the Kripke gluing model for domain types and values leading to the soundness theorem.

4.5 Properties of the Kripke Gluing Model

4.5.1 Monotonicity. As mentioned previously, this model is Kripke in the sense that it is stable under weakenings. We can formalize this property as follows:

$$\begin{array}{c}
\frac{\Gamma \vdash M : A \quad a \in \mathcal{U}_i^* \searrow P \searrow El \quad \downarrow^a m \approx \downarrow^a m \in Nf \quad \forall \Delta, \sigma. \Delta \vdash_w \sigma : \Gamma \rightarrow \Delta \vdash M[\sigma] \approx R_{|\Delta|}^{Nf}(\downarrow^a m) : A[\sigma]}{\Gamma \vdash M : A \text{ @ } m \in Nf_i^a} \\
\frac{\Gamma \vdash M : A \quad a \in \mathcal{U}_i^* \searrow P \searrow El \quad \Gamma \vdash A \text{ @ } P \quad w \approx w \in Ne \quad \forall \Delta, \sigma. \Delta \vdash_w \sigma : \Gamma \rightarrow \Delta \vdash M[\sigma] \approx R_{|\Delta|}^{Ne} w : A[\sigma]}{\Gamma \vdash M : A \text{ @ } w \in Ne_i^a} \\
\frac{\Gamma \vdash A : \text{Type}@i \quad a \approx a \in Ty \quad \forall \Delta, \sigma. \Delta \vdash_w \sigma : \Gamma \rightarrow \Delta \vdash A[\sigma] \approx R_{|\Delta|}^{Ty} w : \text{Type}@i}{\Gamma \vdash A \text{ @ } Ty_i^a}
\end{array}$$

Fig. 10. The Definition of Nf_i^a , Ne_i^a , and Ty_i^a

LEMMA 4.9 (MONOTONICITY). *Given $a \in \mathcal{U}_i^* \searrow P \searrow El$,*

- *For types:* $\Gamma \vdash A \text{ @ } P$ and $\Delta \vdash_w \sigma : \Gamma$ imply $\Delta \vdash A[\sigma] \text{ @ } P$,
- *For terms:* $\Gamma \vdash M : A \text{ @ } m \in El$ and $\Delta \vdash_w \sigma : \Gamma$ imply $\Delta \vdash M[\sigma] : A[\sigma] \text{ @ } m \in El$.

4.5.2 Realizability. As in completeness, the realizability of the Kripke gluing model is a key step towards the soundness theorem. To state and prove the realizability of the Kripke gluing model, we first need to introduce gluing versions of Nf , Ne , and Ty , called Nf_i^a , Ne_i^a , and Ty_i^a . Their definitions are given in Fig. 10.

Note that two main differences exist compared to the previous Nf , Ne , and Ty for completeness. Nf , Ne , and Ty for completeness relate the readbacks of two domain terms/types. On the other hand, the Nf_i^a , Ne_i^a , and Ty_i^a for soundness relate the readback of a domain term/type and a syntactic term/type. More importantly, Nf , Ne , and Ty for completeness relate normal forms using the equality on their syntactic forms whereas Nf_i^a , Ne_i^a , and Ty_i^a use typed equivalence to compare them. These changes allow us to prove the soundness theorem, as the theorem states typed equivalence between a term and its normal form.

Now, we can state the following realizability theorem:

THEOREM 4.10 (REALIZABILITY OF GLUING MODEL). *Given $a \in \mathcal{U}_i^* \searrow P \searrow El$,*

- *For domain values:* $\Gamma \vdash M : A \text{ @ } m \in El$ implies $\Gamma \vdash M : A \text{ @ } m \in Nf_i^a$;
- *For neutral domain values:* $\Gamma \vdash M : A \text{ @ } w \in Ne_i^a$ implies $\Gamma \vdash M : A \text{ @ } \uparrow^a w \in El$;
- *For domain types:* $\Gamma \vdash A \text{ @ } P$ implies $\Gamma \vdash A \text{ @ } Ty_i^a$.

PROOF. By induction on $a \in \mathcal{U}_i^* \searrow P \searrow El$. As in the realizability of the PER model, we need a lemma to deal with $\Gamma \vdash M \text{ @ } m \in Nat$, which is proven by induction on $\Gamma \vdash M \text{ @ } m \in Nat$. $\square$

4.5.3 Other Properties of the Kripke Gluing Model. Again, the fundamental theorem for soundness requires the cumulativity of the gluing model as its key property. However, as the Π case of the gluing model uses El in a negative position, we also need a “lowering” lemma that brings a term into a lower universe. The formal statement of the theorem follows:

LEMMA 4.11 (CUMULATIVITY AND LOWERING OF GLUING MODEL). *Given $a \in \mathcal{U}_i^* \searrow P \searrow El$ and $a \in \mathcal{U}_j^* \searrow P' \searrow El'$ for $i \leq j$,*

- *Cumulativity for types:* $\Gamma \vdash A \text{ @ } P$ implies $\Gamma \vdash A \text{ @ } P'$;
- *Cumulativity for terms:* $\Gamma \vdash M : A \text{ @ } m \in El$ implies $\Gamma \vdash M : A \text{ @ } m \in El'$;
- *Lowering for terms:* $\Gamma \vdash A \text{ @ } P$ and $\Gamma \vdash M : A \text{ @ } m \in El'$ implies $\Gamma \vdash M : A \text{ @ } m \in El$.

We also need a lemma about interaction between the semantic subtyping and the gluing model. First, if two domain types are in the semantic subtyping relation, then glued syntactic types are also in subtyping relation:

LEMMA 4.12 ([SUBTYPING OF GLUING MODEL FOR TYPES](#)). *If $a <_i a'$, $a \in \mathcal{U}_i^* \searrow P \searrow El$, and $a' \in \mathcal{U}_i^* \searrow P' \searrow El'$ then $\Gamma \vdash A \textcircled{R} P$ and $\Gamma \vdash A' \textcircled{R} P'$ implies $\Gamma \vdash A \subseteq A'$.*

Also, the term predicates from the gluing model respect subtyping relation.

LEMMA 4.13 ([SUBTYPING OF GLUING MODEL FOR TERMS](#)). *If $a <_i a'$, $a \in \mathcal{U}_i^* \searrow P \searrow El$, and $a' \in \mathcal{U}_i^* \searrow P' \searrow El'$ then $\Gamma \vdash A' \textcircled{R} P'$ and $\Gamma \vdash M : A \textcircled{R} m \in El$ implies $\Gamma \vdash M : A' \textcircled{R} m \in El'$.*

4.6 Soundness

As for the completeness, we first define the semantic typing judgement $\Gamma \Vdash M : A$ for soundness:

$$\frac{\forall \Delta, \sigma, \rho, P, El. \Delta \vdash \sigma \textcircled{R} \rho \in Sb \rightarrow \llbracket A \rrbracket(\rho) \in \mathcal{U}_i^* \searrow P \searrow El \wedge \Delta \vdash M[\sigma] : A[\sigma] \textcircled{R} \llbracket M \rrbracket(\rho) \in El}{\Gamma \Vdash M : A}$$

Then, we show the following fundamental theorem of the semantic typing judgements for soundness:

THEOREM 4.14 ([FUNDAMENTAL THEOREM FOR SOUNDNESS](#)). *If $\Gamma \vdash M : A$ then $\Gamma \Vdash M : A$.*

PROOF. By induction on the given typing derivation. As in the case of the completeness, we prove each case separately due to its complexity. $\square$

We can derive the soundness of NbE from this:

THEOREM 4.15 ([SOUNDNESS OF NbE](#)). *If $\Gamma \vdash M : A$ then $\Gamma \vdash M \approx nbe_{\Gamma}^A(M') : A$*

PROOF. We first show a lemma that $\Gamma \searrow Sb$ implies $\Gamma \vdash \text{id} \textcircled{R} \uparrow^{\Gamma} \in Sb$ by induction on the given derivation. Then, we prove the main goal by applying the fundamental theorem and this lemma. $\square$

4.7 Key Corollaries

The fundamental theorems for the PER and gluing model allows us to derive several semantic results. Other than the completeness and soundness of NbE, another key result is the consistency of McTT:

COROLLARY 4.16 ([CONSISTENCY](#)). *There is no such M and i satisfying $\Gamma \vdash M : \Pi(x : \text{Type}@i).x$.*

McTT also contains other interesting results such as the idempotency of NbE and injectivity of Π with regard to our equivalence. One can find them in [our semantic consequence module](#).

5 Algorithmic Component

In the previous sections, we have introduced the NbE algorithm that computes $\beta\eta$ normal forms of well-formed expressions, and proved its completeness and soundness. Based on the NbE algorithm, we can provide a set of algorithmic rules for type-checking. For simplicity, our type-checking algorithm is essentially uni-directional; this is possible because we made sure that types can always be inferred from given expressions. During type-checking, we sometimes need to check whether an expression of an inferred type can also have a desired type. For example, in Sec. 2.1, we give an example program where we infer the type $\Pi(y:\text{Type}@0).\text{Type}@0$, but the desired type is $\Pi(y:\text{Type}@0).\text{Type}@1$. Therefore, we also need an algorithm for subtyping. As a result, we are motivated to define the following four algorithmic judgements:




Fig. 11. Algorithmic judgements

- $\Gamma \vdash_A M \Leftarrow A$: Expression M checks against type A .
- $\Gamma \vdash_A M \Rightarrow W$: A *normal* type W is inferred from expression M .
- $\vdash_A^{\text{Nf}} W \subseteq W'$: Normal type W is an algorithmic subtype of normal type W' .
- $\Gamma \vdash_A A \subseteq B$: Type A is an algorithmic subtype of type B .

Selected rules are shown in Fig. 11, and rest of the rules are available on Appendix A (also in [its mechanization](#)). The judgement $\Gamma \vdash_A M \Leftarrow A$ implements type-checking by checking M against a desired type A and is the entry point of the overall type-checking procedure. This judgement first infers a normal type W from M , and then checks if W is a subtype of A .

The main workhorse of the type-checking algorithm is the inference judgement $\Gamma \vdash_A M \Rightarrow W$. Type inference proceeds in a standard manner by recursively walking down the AST so that for any expression, at most one inference rule applies. Note that the judgement always infers a *normal* type, which ensures a unique return for a given input expression. The NbE algorithm is properly invoked to ensure that this invariant is maintained. In certain rules, the checking direction might also be used. For example, in the case of the function application $M N$, we first make sure that M has a function type $\Pi(x : W).W'$. To ensure that the function application is well-formed, $\Gamma \vdash_A N \Leftarrow W'$ is invoked to check whether N has type W' . If that is the case, we can safely say that the overall type of $M N$ is $W'[\text{id}, N/x]$, and we normalize this type before returning it.

Another critical piece that type-checking relies on is the algorithmic subtyping. The main entry is $\Gamma \vdash_A A \subseteq B$, which has only one case. It simply invokes NbE to normalize the input types A and B , and passes the results in normal forms to the core subtyping procedure $\vdash_A^{\text{Nf}} W \subseteq W'$. Note that the context Γ is only used for normalization; the core subtyping comparison algorithm does not rely on the context at all. Fig. 11 shows the full set of three rules for $\vdash_A^{\text{Nf}} W \subseteq W'$. If W and W' are

neither universes nor Π types, then they must be equal. For two universes, the first universe level must be smaller than or equal to the other. Given two Π types, the input types must be the same, while the output type of the first Π must be a subtype of the second output type.

We prove that the algorithmic definitions are sound and complete w.r.t. the declarative definitions. We proceed by first reasoning about the algorithmic subtyping. The easier direction is the soundness of the two algorithmic subtyping judgements:

THEOREM 5.1 (SOUNDNESS OF ALGORITHMIC SUBTYPING FOR NORMAL TYPES). *If $\vdash_A^{Nf} W \subseteq W'$ and for any Γ and i , such that $\Gamma \vdash W : \text{Type}@i$ and $\Gamma \vdash W' : \text{Type}@i$, then $\Gamma \vdash W \subseteq W'$.*

In other words, algorithmic subtyping for normal types always holds regardless of the typing context. Then, the soundness of algorithmic subtyping for general types is proved using the soundness of NbE:

THEOREM 5.2 (SOUNDNESS OF ALGORITHMIC SUBTYPING). *If $\Gamma \vdash_A A \subseteq B$, $\Gamma \vdash A : \text{Type}@i$ and $\Gamma \vdash B : \text{Type}@i$, then $\Gamma \vdash A \subseteq B$.*

The completeness direction is a bit more challenging because we must show that algorithmic subtyping is transitive and works for the Π case.

LEMMA 5.3 (TRANSITIVITY OF ALGORITHMIC SUBTYPING). *If $\Gamma \vdash_A A_0 \subseteq A_1$ and $\Gamma \vdash_A A_1 \subseteq A_2$, then $\Gamma \vdash_A A_0 \subseteq A_2$.*

THEOREM 5.4 (COMPLETENESS OF ALGORITHMIC SUBTYPING). *If $\Gamma \vdash A \subseteq B$, then $\Gamma \vdash_A A \subseteq B$.*

PROOF. Induction on $\Gamma \vdash A \subseteq B$. In the Π case, since we support only covariant subtyping, we know the input types are equivalent. This is checked by the equality of their normal forms. Moreover, the equivalence between input types also ensures that the contexts in which the output types are normalized are also equivalent. Therefore, the normal forms of output types can also be compared by algorithmic subtyping. At this point, the induction hypothesis applies and discharges this case. $\square$

Next, we show that algorithmic typing is also sound and complete w.r.t. the declarative typing. The easier direction is again soundness, which is proved by mutual induction on type-checking and type inference:

THEOREM 5.5 (SOUNDNESS OF ALGORITHMIC TYPING).

- If $\Gamma \vdash_A M \Leftarrow A$, $\vdash \Gamma$ and $\Gamma \vdash A : \text{Type}@i$, then $\Gamma \vdash M : A$.
- If $\Gamma \vdash_A M \Rightarrow W$ and $\vdash \Gamma$, then $\Gamma \vdash M : W$.

Completeness is a collective consequence of all previous results:

THEOREM 5.6 (COMPLETENESS OF ALGORITHMIC TYPING). *If M contains no closure and $\Gamma \vdash M : A$, then $\Gamma \vdash_A M \Leftarrow A$ and there exists W such that $\Gamma \vdash_A M \Rightarrow W$ and $\Gamma \vdash W \subseteq A$.*

6 Implementation and Extraction

So far, we have provided complete and sound NbE algorithm, subtyping algorithm and type-checking algorithm. However, these algorithms cannot be directly implemented as functions because RocQ requires functions to be *total*; while NbE, the base algorithm of all, is partial in general and only terminates when inputs are well-formed. For this reason, in Sec. 3.3 and 5, we define evaluation, readback, and algorithmic typing as functional relations to prove properties about them.

However, one immediate problem is that these functional relations cannot be extracted to executable functions while our purpose is to provide a verified *implementation*. Moreover, an

executable version of these algorithms should not carry proof information during runtime for *efficiency*. Last, the algorithms should be extracted into *human-readable* and comparable to what a skilled human programmer would write. This means that the code can be verified by a skilled human programmer independently. This not only increases our confidence in the final result but also allows developers to experiment with new features before fully verifying them.

To summarize, the main challenge in providing a verified proof checker is to ensure that the extracted code: 1) satisfies the properties proven in the theoretical component; 2) does not carry any extraneous information during runtime; 3) is human-readable. Achieving these three goals at the same time is a nontrivial problem. For example, [Hu et al. \[2023\]](#)'s extraction from Agda to Haskell carries proof information. As a consequence, performance significantly declines and the extracted code base is messy, hard to work with, and difficult to read for a human programmer.

One potential approach to resolve this issue in Rocq is to use the `coq-partialfun` library [\[Winterhalter 2023\]](#), which uses free monads to represent partial functions, so that the latter can be executed within Rocq. This library is used by [Adjedj et al. \[2024\]](#) to implement a type-checking algorithm in Rocq. However, extraction from this library carries some line noise due to the monadic style and is not very human-readable.

In McTT, we use a variation of [Bove \[2009\]](#)'s method, where for each function, we define a call graph to restrict its domain. This call graph also serves as the termination measure. We define this call graph as an inductive relation in Prop. Since Rocq's extraction mechanism omits all instances of Prop, we automatically obtain a code base that is free of proofs and is as readable as a human-written source program. For example, we define the call graph of `eval_exp` (i.e. the name of the functional relation that formalizes $\llbracket M \rrbracket(\rho) = m$ in McTT) as the following inductive type `eval_exp_order`. We only show the application case below:

```
Inductive eval_exp_order : exp -> env -> Prop :=
(* ... *)
| eeo_app : forall M N p,
  eval_exp_order M p ->
  eval_exp_order N p ->
  (forall m n, eval_exp M p m -> eval_exp N p n -> eval_app_order m n) ->
  eval_exp_order (a_app M N) p
```

This case says that the call graph for the function application `a_app M N` case includes the sub-graphs for `M` and `N`, and also another sub-graph for the result of applying the evaluation of `M` to that of `N`.

We now briefly discuss the evaluation function `eval_exp_impl`. As a naming convention, we use the post-fix `impl` to refer to the implementation of the corresponding functional relation `eval_exp`. In addition to the expressions `m:exp` and the environment `p:env`, it also takes in the call graph defined by the inductive type `eval_exp_order m p`. Its return states that there exists a normal form `d` s.t. `eval_exp m p d`. In Rocq, we write the Σ type `{ d | eval_exp m p d }` for it. Mimicking the definition of `eval_exp_order` (i.e. the definition of the call graph), we define `eval_exp_impl` as follows: we first recursively evaluate `M` and `N` using `eval_exp_impl`. Then, we apply `eval_app_impl` (which corresponds to `_·_`) to the results. The function is implemented using the Equations plugin [\[Sozeau and Mangin 2019\]](#) to discharge proof obligations more easily.

```
Equations eval_exp_impl m p (H : eval_exp_order m p) : { d | eval_exp m p d } :=
(* ... *)
| a_app M N , p, H =>
  let (m , Hm) := eval_exp_impl M p _ in
  let (n , Hn) := eval_exp_impl N p _ in
  let (a, Ha) := eval_app_impl m n _ in
  exist _ a _
```

RocQ's extraction removes the call graph and the witness that the resulting normal form d is identical to the one the declarative algorithm returns. In other words, we remove $\text{eval_exp_order } m \ p$ from the input and $\text{eval_exp } m \ p \ d$ from the output. As a result, the extracted code is without any redundant proof witness and is what a skilled human programmer would have written directly in OCaml.

```
let rec eval_exp_impl m p : domain =
(* ... *)
| A_app (e, e0) -> eval_app_impl (eval_exp_impl e p) (eval_exp_impl e0 p)
```

Now, how can we obtain an instance of eval_exp_order when we use this function? Our solution is to prove that eval_exp_order is sound: If there exists a d , s.t. $\text{eval_exp } M \ p \ d$ then we have $\text{eval_exp_order } M \ p$. Our completeness theorem for NbE guarantees that there is a proof of exists d , $\text{eval_exp } M \ p \ d$ for a well-typed term M . In fact, we can easily get this order when we implement a type-checker by type-checking M first.

We have addressed two goals out of three in the implementation: the extracted code carries no extra information and is human-readable. The last problem is to prove that the implementation is indeed sound and complete w.r.t. the corresponding functional relation. In fact, soundness is an immediate consequence of the return type of the function: eval_exp_impl returns $\{ d \mid \text{eval_exp } m \ p \ d \}$, i.e. a domain value d that satisfies $\text{eval_exp } m \ p \ d$. Completeness, on the other hand, says that every instance of a functional relation finds an execution in the implementation. For evaluation, the lemma states that if $\text{eval_exp } m \ p \ d$, then $\text{eval_exp_impl } m \ p$ terminates and returns the same d . This property is easy to conclude, because $\text{eval_exp } m \ p \ d$ is deterministic and eval_exp_order is sound.

We apply this procedure for all algorithms (readback, NbE, subtyping and type-checking) to obtain a fully verified pipeline after extraction satisfying all three criteria.

7 Related Work

7.1 Mechanization of Type Theories

The mechanization of type theories dates back to [Barras and Werner \[1997\]](#)'s effort on mechanizing the Calculus of Constructions. More recently, [Abel et al. \[2018\]](#) mechanize a weak normalization proof for MLTT with two universes using Tait's reducibility candidates [[Tait 1967](#)] in Agda. This mechanization is the basis for many subsequent developments [[Abel et al. 2023](#); [Liu et al. 2025](#); [Pujet and Tabareau 2022, 2023](#)]. Recently, [Adjedj et al. \[2024\]](#) also mechanize this proof in RocQ concentrating on a fixed number of universes. However, Tait-style normalization proofs often require significant technical set-up and time investment when we formalize them in a proof assistant. Furthermore, this style of proof does not directly give us a complete $\beta\eta$ -normalization algorithm.

Our McTT infrastructure uses normalization-by-evaluation (NbE), which often leads to more compact proofs. Normalization by evaluation is a technique to piggyback the normalization problem of the target system onto some computational domain. NbE was originally developed by [Martin-Löf \[1975\]](#) to prove the β normalization of a version of his type theories, and independently by [Berger and Schwichtenberg \[1991\]](#) for their proof checker MinLog. There are different possible choices for the computational domain. One popular choice is based on category theory [[Altenkirch et al. 1995](#)]. Such models are often organized as presheaf categories [[Hu and Pientka 2022](#); [Valliappan et al. 2022](#)], which map the category of weakenings to the set of well-formed terms. [Altenkirch and Kaposi \[2016a,b, 2017\]](#) mechanize an NbE proof for a dependent type theory based on category with families [[Dybjer 1995](#)], a presheaf framework for modelling dependent types.

NbE algorithms are usually easier to understand, extend, and easier to implement if we choose an untyped domain model over a presheaf model [[Abel 2013](#)]. For example, [Abel et al. \[2017\]](#) and [Gratzer et al. \[2019\]](#) both extend this style of NbE proof to sized types and idempotent modal types

respectively. Most recently, [Hu et al. \[2023\]](#) mechanize an NbE algorithm for MLTT extended with a $\Box$ modality, based on [Abel \[2013\]](#)'s untyped domain model in just 11K lines of Agda code. This demonstrates the compactness and elegance of mechanizing NbE for dependent type theories, in particular compared to mechanizations of other styles of normalization proofs.

Mechanizing normalization proofs for type theory within Rocq poses its own challenge: in particular, we cannot rely on induction-recursion in the definition of our semantic model which is common in on-paper proofs as well as in Agda mechanizations. Existing mechanizations of type theory such as by [Adjedj et al. \[2024\]](#) or [Wieczorek and Biernacki \[2018\]](#) cap the number of universes. However, removing this restriction is not straightforward in Rocq. Our McTT infrastructure leverages impredicativity in Rocq together with our subtyping rule to provide a general principled way of developing proofs that usually require induction-recursion. This allows us to mechanize a larger fragment of MLTT than previously in Rocq. Choosing Rocq as our proof assistant also has a distinct advantage: we can leverage Rocq's excellent support for code extraction to derive an OCaml implementation of the verified normalization algorithms without extraneous noises from proof evidences.

7.2 Verified Proof Checking Kernels

In addition to the mechanization of CIC, [Barras and Werner \[1997\]](#) also made the first notable attempt to verify the kernel of a significant fragment of the Calculus of Inductive Construction and to extract a verified type-checker from the verification. More recently, an alternative approach has been pursued in the MetaRocq project to verify the formal semantics of Rocq's type theory within Rocq [\[Sozeau et al. 2020a,b, 2019\]](#). In this line of work, the representation of Rocq terms in the implementation of Rocq is reflected into Rocq itself. Many properties about Rocq can then be proven and mechanized within Rocq about these representations. Since MetaRocq aims to capture the whole feature sets of Rocq, we cannot hope to prove the consistency of Rocq itself due to Gödel's incompleteness theorem. Hence, it is necessary to postulate normalization. The goal of McTT is to actually internalize the normalization proof for a type theory within Rocq. In this approach we will ultimately hit the normalization barrier, but we can still mechanize a significant, but strict subset of Rocq.

Compared to previous work, in McTT, we verify the meta-theory of MLTT hand-in-hand with the algorithmic implementation of a type checker and a normalizer for MLTT. Using Rocq's code extraction, we are able to synthesize an OCaml implementation that is readable and compact. This gives us a verified type-checker for a core fragment of MLTT.

8 Conclusion

We have described the McTT infrastructure to build a verified implementation of a normalizer and type-checker for core MLTT. Every step in the McTT pipeline is verified except for the lexer and pretty-printer. As a result, we have a fully verified kernel for core MLTT with the cumulative universe hierarchy.

In building McTT, we pushed our understanding on how to mechanize type theory in Rocq showing how to leverage impredicativity to deal with the more common induction-recursion definitions. Another important lesson we can draw from our work is that theory and verifiable algorithmic implementations of that theory should be developed hand in hand in a tight loop. In our case, this leads us to include universe subtyping.

Our McTT infrastructure also provides valuable lessons on how to arrive at a verified normalization algorithm within Rocq that allows us to extract readable OCaml code without extraneous noises due to proof witnesses. In particular, we identified two important steps: 1) Encoding partial

functions for evaluation and readback using functional relations and 2) Defining a call graph of each function to restrict the domain of the function following Bove [2009].

We regard McTT as a test-bed for two further directions: First, the theoretical component is a platform for formally verifying standard features like equality types and inductive types, to narrow the gap between the meta-theory and the implementation. Extending our meta-theoretic proofs follows the same principle that we have set up. While this is somewhat tedious, we believe that it is very feasible given the modularity of the McTT infrastructure. This would allow us to extract a verified type-checker to OCaml which can then be used in practice to certify significant parts of existing mechanizations.

Second, McTT can be used to study novel extensions to standard MLTT such as modalities or proof irrelevance. We believe that our theoretical component is sufficiently modular that others can build on it. Furthermore, we believe that leveraging impredicativity for modelling induction-recursion is a technique that is very relevant when mechanizing advanced type theories. More importantly, our McTT infrastructure not only allows us to study the meta-theory of these advanced type theories, but also provides a convenient way to extract a verified type-checker to OCaml, so that we can experiment with these advanced features and explore their potential impact in practice.

Acknowledgments

We would like to thank the anonymous reviewers for their suggestions and feedbacks. This work was funded in part by a Natural Sciences and Engineering Research Council of Canada (discovery grant 206263) and Fonds de recherche du Québec - Nature et technologies (grant 253521). Junyoung Jang was funded by Fonds de recherche du Québec - Nature et technologies (grant 333531). Jason Z. S. Hu was funded partly by Postgraduate Scholarship - Doctoral from the Natural Sciences and Engineering Research Council of Canada and partly by Doctoral (B2X) Research Scholarship from Fonds de recherche du Québec - Nature et technologies during his Ph.D. study.

Data-Availability Statement

This paper is accompanied by a software artifact [Jang et al. 2025] that embodies the key contribution of this paper, a verified kernel of MLTT. The installation instruction and browsable documentation are available in <https://beluga-lang.github.io/McTT/icfp25/>.

References

- Andreas Abel. 2013. *Normalization by Evaluation: Dependent Types and Impredicativity*. Habilitation Thesis. Ludwig-Maximilians-Universität München, Munich, Germany. <https://www.cse.chalmers.se/~abela/habil.pdf>
- Andreas Abel, Nils Anders Danielsson, and Oskar Eriksson. 2023. A Graded Modal Dependent Type Theory with a Universe and Erasure, Formalized. *Proc. ACM Program. Lang.* 7, ICFP (2023), 920–954. doi:10.1145/3607862
- Andreas Abel, Joakim Öhman, and Andrea Vezzosi. 2018. Decidability of Conversion for Type Theory in Type Theory. *Proc. ACM Program. Lang.* 2, POPL (2018), 23:1–23:29. doi:10.1145/3158111
- Andreas Abel, Andrea Vezzosi, and Théo Winterhalter. 2017. Normalization by Evaluation for Sized Dependent Types. *Proc. ACM Program. Lang.* 1, ICFP (2017), 33:1–33:30. doi:10.1145/3110277
- Arthur Adjedi, Meven Lennon-Bertrand, Kenji Maillard, Pierre-Marie Pédro, and Loïc Pujet. 2024. Martin-Löf à la Coq. In *Proceedings of the 13th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2024, London, UK, January 15-16, 2024*, Amin Timany, Dmitriy Traytel, Brigitte Pientka, and Sandrine Blazy (Eds.). ACM, 230–245. doi:10.1145/3636501.3636951
- Thorsten Altenkirch, Martin Hofmann, and Thomas Streicher. 1995. Categorical Reconstruction of a Reduction Free Normalization Proof. In *Proceedings of the 6th International Conference on Category Theory and Computer Science, CTCS 1995, Cambridge, UK, August 7-11, 1995 (Lecture Notes in Computer Science, Vol. 953)*, David H. Pitt, David E. Rydeheard, and Peter T. Johnstone (Eds.). Springer, 182–199. doi:10.1007/3-540-60164-3_27
- Thorsten Altenkirch and Ambrus Kaposi. 2016a. Normalisation by Evaluation for Dependent Types. In *1st International Conference on Formal Structures for Computation and Deduction, FSCD 2016, Porto, Portugal, June 22-26, 2016 (LIPIcs,*

- Vol. 52), Delia Kesner and Brigitte Pientka (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 6:1–6:16. doi:10.4230/LIPICS.FSCD.2016.6
- Thorsten Altenkirch and Ambrus Kaposi. 2016b. Type Theory in Type Theory Using Quotient Inductive Types. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, Florida, USA, January 20-22, 2016*, Rastislav Bodik and Rupak Majumdar (Eds.). ACM, 18–29. doi:10.1145/2837614.2837638
- Thorsten Altenkirch and Ambrus Kaposi. 2017. Normalisation by Evaluation for Type Theory, in Type Theory. *Log. Methods Comput. Sci.* 13, 4 (2017). doi:10.23638/LMCS-13(4:1)2017
- Ali Assaf, Gilles Dowek, Jean-Pierre Jouannaud, and Jiaxiang Liu. 2016. Encoding Proofs in Dedukti: the case of Coq proofs. In *Proceedings Hammers for Type Theories*. 1–6. <https://inria.hal.science/hal-01330980>
- Jeremy Avigad and John Harrison. 2014. Formally verified mathematics. *Commun. ACM* 57, 4 (2014), 66–75.
- Bruno Barras and Benjamin Werner. 1997. Coq in Coq. <https://www.lix.polytechnique.fr/Labo/Bruno.Barras/publi/coqincoq.pdf> Unpublished manuscript.
- Ulrich Berger and Helmut Schwichtenberg. 1991. An Inverse of the Evaluation Functional for Typed Lambda-calculus. In *Proceedings of the 6th Annual Symposium on Logic in Computer Science, LICS 1991, Amsterdam, the Netherlands, July 15-18, 1991*. IEEE Computer Society, 203–211. doi:10.1109/LICS.1991.151645
- Mathieu Boespflug and Guillaume Burel. 2012. CoqInE: Translating the Calculus of Inductive Constructions into the $\lambda\Pi$ -calculus Modulo. In *Workshop on ProofExchange for Theorem Proving (PxTP) (CEUR Workshop Proceedings, Vol. 878)*. CEUR-WS.org, 44–50.
- Ana Bove. 2009. Another Look at Function Domains. *Electronic Notes in Theoretical Computer Science* 249 (2009), 61–74. doi:10.1016/j.entcs.2009.07.084 Proceedings of the 25th Conference on Mathematical Foundations of Programming Semantics (MFPS 2009).
- Ana Bove and Venanzio Capretta. 2005. Modelling general recursion in type theory. *Mathematical Structures in Computer Science* 15, 4 (2005), 671–708. doi:10.1017/S0960129505004822
- Kevin Buzzard, Johan Commelin, and Patrick Massot. 2020. Formalising perfectoid spaces. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (New Orleans, LA, USA) (CPP 2020)*. Association for Computing Machinery, New York, NY, USA, 299–312. doi:10.1145/3372885.3373830
- James Chapman. 2008. Type Theory Should Eat Itself. In *Proceedings of the International Workshop on Logical Frameworks and Metalanguages: Theory and Practice, LFMT@LICS 2008, Pittsburgh, Pennsylvania, USA, June 23, 2008 (Electronic Notes in Theoretical Computer Science, Vol. 228)*, Andreas Abel and Christian Urban (Eds.). Elsevier, 21–36. doi:10.1016/J.ENTCS.2008.12.114
- Thierry Coquand. 2018. Canonicity and normalisation for Dependent Type Theory. CoRR abs/1810.09367 (2018). arXiv:1810.09367 <http://arxiv.org/abs/1810.09367>
- Thierry Coquand and Peter Dybjer. 1997. Intuitionistic model constructions and normalization proofs. *Mathematical Structures in Computer Science* 7, 1 (1997), 75–94. doi:10.1017/S0960129596002150
- Thierry Coquand and Christine Paulin. 1988. Inductively Defined Types. In *COLOG-88, International Conference on Computer Logic, Tallinn, USSR, December 1988, Proceedings (Lecture Notes in Computer Science, Vol. 417)*, Per Martin-Löf and Grigori Mints (Eds.). Springer, 50–66. doi:10.1007/3-540-52335-9_47
- Peter Dybjer. 1995. Internal Type Theory. In *International Workshop on Types for Proofs and Programs, TYPES 1995, Torino, Italy, June 5-8, 1995, Selected Papers (Lecture Notes in Computer Science, Vol. 1158)*, Stefano Berardi and Mario Coppo (Eds.). Springer, 120–134. doi:10.1007/3-540-61780-9_66
- Peter Dybjer and Anton Setzer. 2003. Induction-recursion and Initial Algebras. *Ann. Pure Appl. Log.* 124, 1-3 (2003), 1–47. doi:10.1016/S0168-0072(02)00096-9
- Daniel Gratzer, Jonathan Sterling, and Lars Birkedal. 2019. Implementing A Modal Dependent Type Theory. *Proc. ACM Program. Lang.* 3, ICFP (2019), 107:1–107:29. doi:10.1145/3341711
- Robert Harper and Frank Pfenning. 2005. On equivalence and canonical forms in the LF type theory. *ACM Trans. Comput. Log.* 6, 1 (2005), 61–101. doi:10.1145/1042038.1042041
- Jason Z. S. Hu, Junyoung Jang, and Brigitte Pientka. 2023. Normalization by Evaluation for Modal Dependent Type Theory. *J. Funct. Program.* 33 (2023). doi:10.1017/S0956796823000060
- Jason Z. S. Hu and Brigitte Pientka. 2022. A Categorical Normalization Proof for the Modal Lambda-Calculus. In *Proceedings of the 38th Conference on the Mathematical Foundations of Programming Semantics, MFPS 2022, Cornell University, Ithaca, New York, USA, with a satellite event at IRIF, Denis Diderot University, Paris, France, and online, July 11-13, 2022 (EPTICS, Vol. 1)*, Justin Hsu and Christine Tasson (Eds.). EpiSciences. doi:10.46298/ENTICS.10360
- Junyoung Jang, Antoine Gaulin, Jason Z. S. Hu, and Brigitte Pientka. 2025. McTT: A Verified Kernel for a Proof Assistant. doi:10.5281/zenodo.15712175
- Dominique Larchey-Wendling and Jean-François Monin. 2018. Simulating Induction-Recursion for Partial Algorithms. In *24th International Conference on Types for Proofs and Programs, TYPES 2018, Braga, Portugal*. <https://hal.science/hal-02333374>
- Xavier Leroy. 2009a. Formal verification of a realistic compiler. *Commun. ACM* 52, 7 (2009), 107–115.

- Xavier Leroy. 2009b. A Formally Verified Compiler Back-end. *J. Autom. Reasoning* 43, 4 (2009), 363–446.
- Yiyun Liu, Jonathan Chan, and Stephanie Weirich. 2025. Consistency of a Dependent Calculus of Indistinguishability. *Proc. ACM Program. Lang.* 9, POPL, Article 7 (Jan. 2025), 27 pages. doi:10.1145/3704843
- Per Martin-Löf. 1984. *Intuitionistic Type Theory*. Studies in proof theory, Vol. 1. Bibliopolis.
- Per Martin-Löf. 1975. An Intuitionistic Theory of Types: Predicative Part. In *Logic Colloquium 1973*, H.E. Rose and J.C. Shepherdson (Eds.). Studies in Logic and the Foundations of Mathematics, Vol. 80. Elsevier, 73–118. doi:10.1016/S0049-237X(08)71945-1
- Christine Paulin-Mohring. 1993. Inductive Definitions in the System Coq - Rules and Properties. In *Typed Lambda Calculi and Applications, International Conference on Typed Lambda Calculi and Applications, TLCA '93, Utrecht, the Netherlands, March 16-18, 1993, Proceedings (Lecture Notes in Computer Science, Vol. 664)*, Marc Bezem and Jan Friso Groote (Eds.). Springer, 328–345. doi:10.1007/BFB0037116
- Loïc Pujet and Nicolas Tabareau. 2022. Observational Equality: Now for Good. *Proc. ACM Program. Lang.* 6, POPL (2022), 1–27. doi:10.1145/3498693
- Loïc Pujet and Nicolas Tabareau. 2023. Impredicative Observational Equality. *Proc. ACM Program. Lang.* 7, POPL (2023), 2171–2196. doi:10.1145/3571739
- John C. Reynolds. 1998. Definitional Interpreters for Higher-Order Programming Languages. *High. Order Symb. Comput.* 11, 4 (1998), 363–397. doi:10.1023/A:1010027404223
- Matthieu Sozeau, Abhishek Anand, Simon Boulrier, Cyril Cohen, Yannick Forster, Fabian Kunze, Gregory Malecha, Nicolas Tabareau, and Théo Winterhalter. 2020a. The MetaCoq Project. *J. Autom. Reason.* 64, 5 (2020), 947–999. https://doi.org/10.1007/s10817-019-09540-0
- Matthieu Sozeau, Abhishek Anand, Simon Boulrier, Cyril Cohen, Yannick Forster, Fabian Kunze, Gregory Malecha, Nicolas Tabareau, and Théo Winterhalter. 2020b. The MetaCoq Project. *Journal of Automated Reasoning* (Feb. 2020). doi:10.1007/s10817-019-09540-0
- Matthieu Sozeau, Simon Boulrier, Yannick Forster, Nicolas Tabareau, and Théo Winterhalter. 2019. Coq Coq correct! verification of type checking and erasure for Coq, in Coq. *Proc. ACM Program. Lang.* 4, POPL, Article 8 (Dec. 2019), 28 pages. doi:10.1145/3371076
- Matthieu Sozeau and Cyprien Mangin. 2019. Equations reloaded: high-level dependently-typed functional programming and proving in Coq. *Proc. ACM Program. Lang.* 3, ICFP (2019), 86:1–86:29. doi:10.1145/3341690
- William W. Tait. 1967. Intensional Interpretations of Functionals of Finite Type I. *J. Symb. Log.* 32, 2 (1967), 198–212. doi:10.2307/2271658
- The Mathlib Community. 2020. The Lean Mathematical Library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2020, New Orleans, Louisiana, USA, January 20-21, 2020*, Jasmin Blanchette and Catalin Hritcu (Eds.). ACM, 367–381. doi:10.1145/3372885.3373824
- Nachiappan Valliappan, Fabian Ruch, and Carlos Tomé Cortiñas. 2022. Normalization for Fitch-style Modal Calculi. *Proc. ACM Program. Lang.* 6, ICFP (2022), 772–798. https://doi.org/10.1145/3547649
- Paweł Wieczorek and Dariusz Biernacki. 2018. A Coq Formalization of Normalization by Evaluation for Martin-Löf Type Theory. In *Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2018, Los Angeles, California, USA, January 8-9, 2018*, June Andronick and Amy P. Felty (Eds.). ACM, 266–279. doi:10.1145/3167091
- Théo Winterhalter. 2023. Composable partial functions in Coq, totally for free. In *29th International Conference on Types for Proofs and Programs TYPES 2023 — Abstracts*. 208.

A Definition of Martin-Löf type theory

A.1 Syntactic definition

A.1.1 Syntax. We recall the syntax of McTT for convenience:

Variables			x, y, r
De Bruijn indices	$\mathbb{N}$	$\ni$	d
Universe levels	$\mathbb{N}$	$\ni$	i, j
Contexts	Ctx	$\ni$	$\Gamma, \Delta ::= \cdot \mid \Gamma, x : A$
Expressions	Exp	$\ni$	$A, B, C, M, N ::= x_d \mid \text{Nat} \mid \text{Type}@i \mid \Pi(x:A).B$ $\mid \text{zero} \mid \text{succ } M \mid \text{rec}_{x:A} M_{\text{zero}} (y, r. M_{\text{succ}}) N$ $\mid \lambda(x:A).M \mid M N \mid M[\sigma]$
Substitutions	Subst	$\ni$	$\sigma, \delta ::= \text{id} \mid \text{wk} \mid \sigma \circ \delta \mid \sigma, M/x_0$

A.1.2 Judgments. We give a more exhaustive definition of the declarative judgements of McTT. Congruence rules are omitted from the equivalence judgments.

$\boxed{\vdash \Gamma}$ Γ is a well-formed context.

$$\frac{}{\vdash \cdot} \quad \frac{\vdash \Gamma \quad \Gamma \vdash A : \text{Type}@i}{\vdash \Gamma, x:A}$$

$\boxed{\Gamma \vdash t : T}$ Term t has type T in Γ .

$$\begin{array}{c} \frac{\vdash \Gamma \quad x : A \in \Gamma}{\Gamma \vdash x : A} \quad \frac{\vdash \Gamma \quad \Gamma \vdash A : \text{Type}@i}{\Gamma \vdash A : \text{Type}@1+i} \quad \frac{\vdash \Gamma}{\Gamma \vdash \text{Nat} : \text{Type}@0} \quad \frac{\vdash \Gamma}{\Gamma \vdash \text{zero} : \text{Nat}} \\ \\ \frac{\Gamma \vdash M : \text{Nat}}{\Gamma \vdash \text{succ } M : \text{Nat}} \quad \frac{\Gamma, x : \text{Nat} \vdash A : \text{Type}@i \quad \Gamma \vdash M_{\text{zero}} : A[\text{id}, \text{zero}/x] \quad \Gamma, y : \text{Nat}, r : A \vdash M_{\text{succ}} : A[\text{wk} \circ \text{wk}, \text{succ } y/x] \quad \Gamma \vdash N : \text{Nat}}{\Gamma \vdash \text{rec}_{x:A} M_{\text{zero}} (y, r. M_{\text{succ}}) N : A[\text{id}, N/x]} \\ \\ \frac{\Gamma \vdash A : \text{Type}@i \quad \Gamma, x : A \vdash B : \text{Type}@i}{\Gamma \vdash \Pi(x : A).B : \text{Type}@i} \quad \frac{\Gamma \vdash A : \text{Type}@i \quad \Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda(x : A).M : \Pi(x : A).B} \\ \\ \frac{\Gamma \vdash A : \text{Type}@i \quad \Gamma, x : A \vdash B : \text{Type}@i}{\Gamma \vdash M : \Pi(x : A).B} \quad \frac{\Gamma \vdash N : A}{\Gamma \vdash M N : B[\text{id}, N/x]} \quad \frac{\Delta \vdash M : A \quad \Gamma \vdash \sigma : \Delta}{\Gamma \vdash M[\sigma] : A[\sigma]} \\ \\ \frac{\Gamma \vdash A' : \text{Type}@i \quad \Gamma \vdash M : A \quad \Gamma \vdash A \subseteq A'}{\Gamma \vdash M : A'} \end{array}$$

$\boxed{\Gamma \vdash \sigma : \Delta}$ σ is a well-formed substitution from Δ to Γ .

$$\begin{array}{c} \frac{\vdash \Gamma}{\Gamma \vdash \text{id} : \Gamma} \quad \frac{\vdash \Gamma, x : A}{\Gamma, x : A \vdash \text{wk} : \Gamma} \quad \frac{\Gamma \vdash \sigma : \Delta \quad \Delta \vdash A : \text{Type}@i \quad \Gamma \vdash M : A[\sigma]}{\Gamma \vdash \sigma, M/x_0 : \Delta, x : A} \\ \\ \frac{\Gamma' \vdash \sigma : \Gamma'' \quad \Gamma \vdash \delta : \Gamma'}{\Gamma \vdash \sigma \circ \delta : \Gamma''} \quad \frac{\boxed{\vdash \Delta'} \quad \Gamma \vdash \sigma : \Delta \quad \vdash \Delta \subseteq \Delta'}{\Gamma \vdash \sigma : \Delta'} \end{array}$$

$\boxed{\Gamma \vdash M \approx N : A}$ M and N of type A are equivalent.

$$\begin{array}{c}
\frac{\Gamma, x : \text{Nat} \vdash A : \text{Type}@i \quad \Gamma \vdash M_{\text{zero}} : A[\text{id}, \text{zero}/x_0] \quad \Gamma, y : \text{Nat}, r : A \vdash M_{\text{succ}} : A[\text{wk} \circ \text{wk}, \text{succ } y/x_0]}{\Gamma \vdash \text{rec}_{x:A} M_{\text{zero}} (y, r.M_{\text{succ}}) \text{ zero} \approx M_{\text{zero}} : A[\text{id}, \text{zero}/x_0]} \\
\frac{\Gamma, x : \text{Nat} \vdash A : \text{Type}@i \quad \Gamma \vdash M_{\text{zero}} : A[\text{id}, \text{zero}/x_0] \quad \Gamma, y : \text{Nat}, r : A \vdash M_{\text{succ}} : A[\text{wk} \circ \text{wk}, \text{succ } y/x_0] \quad \Gamma \vdash N : \text{Nat} \quad M_r = \text{rec}_{x:A} M_{\text{zero}} (y, r.M_{\text{succ}}) N}{\Gamma \vdash \text{rec}_{x:A} M_{\text{zero}} (y, r.M_{\text{succ}}) (\text{succ } N) \approx M_{\text{succ}}[\text{id}, N/y_1, M_r/r_0] : A[\text{id}, \text{succ } N/x_0]} \\
\frac{\Gamma \vdash A : \text{Type}@i \quad \Gamma, x : A \vdash B : \text{Type}@i \quad \Gamma, x : A \vdash M : B \quad \Gamma \vdash N : A}{\Gamma \vdash (\lambda(x : A).M) N \approx M[\text{id}, N/x_0] : B[\text{id}, N/x_0]} \\
\frac{\Gamma \vdash A : \text{Type}@i \quad \Gamma, x : A \vdash B : \text{Type}@i \quad \Gamma \vdash M : \Pi(x : A).B}{\Gamma \vdash M \approx \lambda(x : A).(M[\text{wk}]) x : \Pi(x : A).B} \\
\frac{\Gamma \vdash A' : \text{Type}@i \quad \Gamma \vdash M \approx N : A \quad \Gamma \vdash A \subseteq A'}{\Gamma \vdash M \approx N : A'}
\end{array}$$

$\boxed{\Gamma \vdash \sigma \approx \delta : \Delta}$ σ and δ are equivalent substitutions from Δ to Γ .

$$\begin{array}{c}
\frac{\Gamma' \vdash \sigma : \Gamma'' \quad \Gamma'' \vdash A : \text{Type}@i \quad \Gamma' \vdash M : A[\sigma] \quad \Gamma \vdash \delta : \Gamma'}{\Gamma \vdash (\sigma, M/x_0) \circ \delta \approx (\sigma \circ \delta), M[\delta]/x_0 : \Gamma'', x : A} \\
\frac{\Gamma \vdash \sigma : \Delta \quad \Delta \vdash A : \text{Type}@i \quad \Gamma \vdash M : A[\sigma]}{\Gamma \vdash \text{wk} \circ (\sigma, M/x_0) \approx \sigma : \Delta} \\
\frac{\Gamma \vdash \sigma \approx \sigma' : \Delta \quad \vdash \Delta \subseteq \Delta'}{\Gamma \vdash \sigma \approx \sigma' : \Delta'} \\
\frac{\Gamma \vdash \sigma : \Gamma', x : A}{\Gamma \vdash \sigma \approx (\text{wk} \circ \sigma), x[\sigma]/x_0 : \Gamma, x : A}
\end{array}$$

$\boxed{\Gamma \vdash A \subseteq B}$ A is a subtype of B .

$$\begin{array}{c}
\frac{\Gamma \vdash B : \text{Type}@i \quad \Gamma \vdash A \subseteq A' \quad \Gamma \vdash A' \subseteq A''}{\Gamma \vdash A \subseteq B} \quad \frac{\vdash \Gamma \quad i < j}{\Gamma \vdash \text{Type}@i \subseteq \text{Type}@j} \\
\frac{\Gamma \vdash A : \text{Type}@i \quad \Gamma \vdash A' : \text{Type}@i \quad \Gamma \vdash A \approx A' : \text{Type}@i \quad \Gamma, x : A \vdash B : \text{Type}@i \quad \Gamma, x : A' \vdash B' : \text{Type}@i \quad \Gamma, x : A' \vdash B \subseteq B'}{\Gamma \vdash \Pi(x : A).B \subseteq \Pi(x : A').B'}
\end{array}$$

$\boxed{\vdash \Gamma \subseteq \Delta}$ Γ is a sub-context of Δ .

$$\frac{\vdash \Gamma \subseteq \Gamma' \quad \Gamma \vdash A : \text{Type}@i \quad \Gamma' \vdash A' : \text{Type}@i \quad \Gamma \vdash A \subseteq A'}{\vdash \Gamma, x : A \subseteq \Gamma', x : A'}$$

A.2 Semantic definition

A.2.1 *Domain.* We recall the definition of the domain for convenience:

De Bruijn levels	$\mathbb{N}$	$\ni$	z	
Neutral domain values	D^{Ne}	$\ni$	v	$::= l_z \mid \text{rec}(A, b, M_{\text{succ}}, \rho, v) \mid v \ w$
Domain values	D	$\ni$	a, b, m, n	$::= \uparrow^a (v) \mid N \mid \text{U@}i \mid \Pi(a, B, \rho)$
				$\mid \text{ze} \mid \text{su}(a) \mid \Lambda(M, \rho)$
Normal domain values	D^{Nf}	$\ni$	w	$::= \downarrow^a (m)$
Evaluation environments	Env	$\ni$	ρ	

A.2.2 *Evaluation.*

$$\begin{aligned}
& \llbracket _ \rrbracket (_) :: \text{Exp} \rightarrow \text{Env} \rightarrow D \\
& \llbracket x_d \rrbracket (\rho) := \rho(d) \\
& \llbracket \Pi(x : A).B \rrbracket (\rho) := \Pi(\llbracket A \rrbracket (\rho), B, \rho) \\
& \llbracket \text{rec}_{x.A} M_{\text{zero}} (y, r.M_{\text{succ}} N) \rrbracket (\rho) := \text{rec} \cdot (A, \llbracket M_{\text{zero}} \rrbracket (\rho), M_{\text{succ}}, \rho, \llbracket N \rrbracket (\rho)) \\
& \llbracket \lambda(x : A).M \rrbracket (\rho) := \Lambda(M, \rho) \\
& \llbracket M N \rrbracket (\rho) := \llbracket M \rrbracket (\rho) \cdot \llbracket N \rrbracket (\rho) \\
& \llbracket t[\sigma] \rrbracket (\rho) := \llbracket t \rrbracket (\llbracket \sigma \rrbracket (\rho)) \\
& \llbracket _ \rrbracket (_) :: \text{Subst} \rightarrow \text{Env} \rightarrow \text{Env} \\
& \llbracket \text{id} \rrbracket (\rho) := \rho \\
& \llbracket \text{wk} \rrbracket (\rho) := \text{drop}(\rho) \\
& \llbracket \sigma, M/x_0 \rrbracket (\rho) := \text{ext}(\llbracket \sigma \rrbracket (\rho), \llbracket M \rrbracket (\rho)) \\
& \llbracket \sigma \circ \delta \rrbracket (\rho) := \llbracket \sigma \rrbracket (\llbracket \delta \rrbracket (\rho)) \\
& _ \cdot _ :: D \rightarrow D \rightarrow D \\
& (\Lambda(M, \rho)) \cdot a := \llbracket M \rrbracket (\text{ext}(\rho, a)) \\
& (\uparrow^{\Pi(a, B, \rho)} (v)) \cdot m := \uparrow^{\llbracket B \rrbracket (\text{ext}(\rho, m))} (v \downarrow^a (m)) \\
& \text{rec} \cdot (_, _, _, _) :: \text{Exp} \rightarrow D \rightarrow \text{Exp} \rightarrow \text{Env} \rightarrow D \rightarrow D \\
& \text{rec} \cdot (A, m_{\text{zero}}, M_{\text{succ}}, \rho, \text{ze}) := m_{\text{zero}} \\
& \text{rec} \cdot (A, m_{\text{zero}}, M_{\text{succ}}, \rho, \text{su}(m)) := \llbracket M_{\text{succ}} \rrbracket (\text{ext}(\rho, m, \text{rec} \cdot (A, m_{\text{zero}}, M_{\text{succ}}, \rho, m))) \\
& \text{rec} \cdot (A, m_{\text{zero}}, M_{\text{succ}}, \rho, \uparrow^b (v)) := \uparrow^{\llbracket A \rrbracket (\text{ext}(\rho, \uparrow^b (v)))} (\text{rec}(A, m_{\text{zero}}, M_{\text{succ}}, \rho, v))
\end{aligned}$$

A.2.3 *Readback.*

$$\begin{aligned}
& R^{\text{Nf}} :: \mathbb{N} \rightarrow D^{\text{Nf}} \rightarrow \text{Nf} \\
& R_z^{\text{Nf}}(\downarrow^a (\uparrow^b (v))) := R_z^{\text{Ne}}(v) \quad (\text{where } a = \uparrow^{a'}(c) \text{ or } N) \\
& R_z^{\text{Nf}}(\downarrow^{\text{U@}i} (a)) := R_z^{\text{Ty}}(a) \\
& R_z^{\text{Nf}}(\downarrow^{\Pi(a, B, \rho)} (m)) := \lambda(x : R_z^{\text{Ty}}(a)). R_{1+z}^{\text{Nf}}(\downarrow^{\llbracket B \rrbracket (\text{ext}(\rho, \uparrow^a(l_z)))} (m \cdot \uparrow^a(l_z))) \\
& R^{\text{Ty}} :: \mathbb{N} \rightarrow D \rightarrow \text{Nf} \\
& R_z^{\text{Ty}}(\Pi(a, B, \rho)) := \Pi(x : R_z^{\text{Ty}}(a)). R_{1+z}^{\text{Ty}}(\llbracket B \rrbracket (\text{ext}(\rho, \uparrow^a(l_z)))) \\
& R_z^{\text{Ty}}(\uparrow^a (v)) := R_z^{\text{Ne}}(v) \\
& R^{\text{Ne}} :: \mathbb{N} \rightarrow D^{\text{Ne}} \rightarrow \text{Ne} \\
& R_z^{\text{Ne}}(l_z) := x_{\max(z' - z - 1, 0)} \\
& R_z^{\text{Ne}}(\text{rec}(A, b, M_{\text{succ}}, \rho, v)) := \text{rec}_{x.W} R_z^{\text{Nf}}(\downarrow^{\llbracket A \rrbracket (\text{ext}(\rho, \text{ze}))} (b)) (x, y, W_{\text{succ}} R_z^{\text{Ne}}(v)) \\
& \quad (\text{where } a := \llbracket A \rrbracket (\text{ext}(\rho, \uparrow^N(l_z))), W := R_{1+z}^{\text{Ty}}(a), \text{ and}) \\
& (a' := \llbracket A \rrbracket (\text{ext}(\rho, \text{su}(\uparrow^N(l_z))))), m := \llbracket M_{\text{succ}} \rrbracket (\text{ext}(\rho, \uparrow^N(l_z), \uparrow^a(l_{1+z}))), W_{\text{succ}} := R_{2+z}^{\text{Nf}}(\downarrow^{a'}(m))) \\
& R_z^{\text{Ne}}(v \ w) := R_z^{\text{Ne}}(v) R_z^{\text{Nf}}(w)
\end{aligned}$$

A.3 Algorithmic judgements

 $\boxed{\vdash_A^{\text{Nf}} W \subseteq W'}$

Algorithmic subtyping between normal types

$$\frac{W \text{ is not a universe or a } \Pi \text{ type}}{\vdash_A^{\text{Nf}} W \subseteq W} \quad \frac{i \leq j}{\vdash_A^{\text{Nf}} \text{Type}@i \subseteq \text{Type}@j}$$

$$\frac{\vdash_A^{\text{Nf}} W_2 \subseteq W'_2}{\vdash_A^{\text{Nf}} \Pi(x : W_1).W_2 \subseteq \Pi(x : W_1).W'_2}$$

 $\boxed{\Gamma \vdash_A A \subseteq B}$

Algorithmic subtyping between well-formed types

$$\frac{\vdash_A^{\text{Nf}} \text{nbe}_\Gamma(A) \subseteq \text{nbe}_\Gamma(B)}{\Gamma \vdash_A A \subseteq B}$$

 $\boxed{\Gamma \vdash_A M \Leftarrow A}$

Term M is checked against type A

$$\frac{\Gamma \vdash_A M \Rightarrow W \quad \Gamma \vdash_A A \subseteq B}{\Gamma \vdash_A M \Leftarrow B}$$

 $\boxed{\Gamma \vdash_A M \Rightarrow W}$

W is an inferred normal type of term M

$$\frac{x : A \in \Gamma}{\Gamma \vdash_A x \Rightarrow \text{nbe}_\Gamma(A)} \quad \frac{}{\Gamma \vdash_A \text{Type}@i \Rightarrow \text{Type}@ (1+i)} \quad \frac{}{\Gamma \vdash_A \text{Nat} \Rightarrow \text{Type}@0}$$

$$\frac{}{\Gamma \vdash_A \text{zero} \Rightarrow \text{Nat}} \quad \frac{\Gamma \vdash_A M \Leftarrow \text{Nat}}{\Gamma \vdash_A \text{succ } M \Rightarrow \text{Nat}}$$

$$\frac{\Gamma, x : \text{Nat} \vdash_A A \Rightarrow \text{Type}@i \quad \Gamma \vdash_A M_{\text{zero}} \Leftarrow A[\text{id}, \text{zero}/x] \quad \Gamma, y : \text{Nat}, r : A \vdash_A M_{\text{succ}} \Leftarrow A[(\text{wk} \circ \text{wk}), \text{succ } y/x] \quad \Gamma \vdash_A N \Leftarrow \text{Nat}}{\Gamma \vdash_A \text{rec}_{x:A} M_{\text{zero}} (y, r.M_{\text{succ}}) N \Rightarrow \text{nbe}_\Gamma(A[\text{id}, N/x])}$$

$$\frac{\Gamma \vdash_A A \Rightarrow \text{Type}@i \quad \Gamma, x : A \vdash_A B \Rightarrow \text{Type}@j}{\Gamma \vdash_A \Pi(x : A).B \Rightarrow \text{Type}@ \max(i, j)} \quad \frac{\Gamma \vdash_A A \Rightarrow \text{Type}@i \quad \Gamma, x : A \vdash_A M \Rightarrow W}{\Gamma \vdash_A \lambda(x : A).M \Rightarrow \Pi(x : \text{nbe}_\Gamma(A)).W}$$

$$\frac{\Gamma \vdash_A M \Rightarrow \Pi(x : W).W' \quad \Gamma \vdash_A N \Leftarrow W}{\Gamma \vdash_A M N \Rightarrow \text{nbe}_\Gamma(W'[\text{id}, N/x])}$$

Received 2025-02-27; accepted 2025-06-27