Approximate Normalization for Gradual Dependent Types

JOSEPH EREMONDI, University of British Columbia, Canada

ÉRIC TANTER, University of Chile, Chile and Inria Paris, France

RONALD GARCIA, University of British Columbia, Canada

Dependent types help programmers write highly reliable code. However, this reliability comes at a cost: it can be challenging to write new prototypes in (or migrate old code to) dependently-typed programming languages. Gradual typing makes static type disciplines more flexible, so an appropriate notion of gradual dependent types could fruitfully lower this cost. However, dependent types raise unique challenges for gradual typing. Dependent typechecking involves the execution of program code, but gradually-typed code can signal runtime type errors or diverge. These runtime errors threaten the soundness guarantees that make dependent types so attractive, while divergence spoils the type-driven programming experience.

This paper presents GDTL, a gradual dependently-typed language that emphasizes pragmatic dependently-typed programming. GDTL fully embeds both an untyped and dependently-typed language, and allows for smooth transitions between the two. In addition to gradual types we introduce *gradual terms*, which allow the user to be imprecise in type indices and to omit proof terms; runtime checks ensure type safety. To account for nontermination and failure, we distinguish between compile-time normalization and run-time execution: compile-time normalization is *approximate* but total, while runtime execution is *exact*, but may fail or diverge. We prove that GDTL has decidable typechecking and satisfies all the expected properties of gradual languages. In particular, GDTL satisfies the static and dynamic gradual guarantees: reducing type precision preserves typedness, and altering type precision does not change program behavior outside of dynamic type failures. To prove these properties, we were led to establish a novel *normalization gradual guarantee* that captures the monotonicity of approximate normalization with respect to imprecision.

CCS Concepts: • **Theory of computation** → **Type structures; Program semantics.**

Additional Key Words and Phrases: Gradual types, dependent types, normalization

ACM Reference Format:

Joseph Eremondi, Éric Tanter, and Ronald Garcia. 2019. Approximate Normalization for Gradual Dependent Types. *Proc. ACM Program. Lang.* 3, ICFP, Article 88 (August 2019), 30 pages. <https://doi.org/10.1145/3341692>

1 INTRODUCTION

Dependent types support the development of extremely reliable software. With the full power of higher-order logic, programmers can write expressive specifications as types, and be confident that if a program typechecks, then it meets its specification. Dependent types are at the core of proof assistants like Coq [Bertot and Castéran 2004] and Agda [Norell 2009]. While these pure systems can be used for certified programming [Chlipala 2013], their focus is on the construction of proofs,

*This work is partially funded by CONICYT FONDECYT Regular Project 1190058, ERC Starting Grant SECOMP (715753), an NSERC Discovery grant, and the NSERC Canada Graduate Scholarship.

Authors' addresses: Joseph Eremondi, Department of Computer Science, University of British Columbia, Canada, jeremond@cs.ubc.ca; Éric Tanter, Computer Science Department (DCC), University of Chile, Chile, Inria Paris, France, etanter@dcc.uchile.cl; Ronald Garcia, Department of Computer Science, University of British Columbia, Canada, rxg@cs.ubc.ca.

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

© 2019 Copyright held by the owner/author(s).

2475-1421/2019/8-ART88

<https://doi.org/10.1145/3341692>

rather than practical or efficient code. Dependently-typed extensions of practical programming languages maintain a clear phase distinction between compile-time typechecking and runtime execution, and have to embrace some compromise regarding impurity. One possibility is to forbid potentially impure expressions from occurring in types, either by considering a separate pure sub-language of type-level computation as in Dependent ML [Xi and Pfenning 1999], by using an effect system and termination checker to prevent impurity to leak in type dependencies as in F★ [Swamy et al. 2016] and Idris [Brady 2013], or by explicitly separating the language into two fragments with controlled interactions between them, as in Zombie [Casinghino et al. 2014; Sjöberg et al. 2012]. A radical alternative is to give up on decidable typechecking and logical consistency altogether and give all responsibility to the programmer, as in Dependent Haskell [Eisenberg 2016]. The design space is wide, and practical dependently-typed programming is a fertile area of research.

As with any static type system, the reliability brought by dependent types comes at a cost. Dependently-typed languages impose a rigid discipline, sometimes requiring programmers to explicitly construct proofs in their programs. Because of this rigidity, dependent types can interfere with rapid prototyping, and migrating code from languages with simpler type systems can be difficult. Several approaches have been proposed to relax dependent typing in order to ease programming, for instance by supporting some form of interoperability between (possibly polymorphic) non-dependently-typed and dependently-typed programs and structures [Dagand et al. 2018; Osera et al. 2012; Ou et al. 2004; Tanter and Tabareau 2015]. These approaches require programmers to explicitly trigger runtime checks through casts, liftings, or block boundaries. Such explicit interventions hamper evolution.

In contrast, gradual typing [Siek and Taha 2006] exploits *type imprecision* to drive the interaction between static and dynamic checking in a smooth, continuous manner [Siek et al. 2015]. A gradual language introduces an unknown type $?$, and admits imprecise types such as $\text{Nat} \rightarrow ?$. The gradual type system optimistically handles imprecision, deferring to runtime checks where needed. Therefore, runtime checking is an *implicit* consequence of type imprecision, and is seamlessly adjusted as programmers evolve the declared types of components, be they modules, functions, or expressions. This paper extends gradual typing to provide a flexible incremental path to adopting dependent types.

Gradual typing has been adapted to many other type disciplines, including ownership types [Sergey and Clarke 2012], effects [Bañados Schwerter et al. 2016], refinement types [Lehmann and Tanter 2017], security types [Fennell and Thiemann 2013; Toro et al. 2018a], and session types [Igarashi et al. 2017]. But it has not yet reached dependent types. Even as the idea holds much promise, it also poses significant challenges. The greatest barrier to gradual dependent types is that a dependent type checker must evaluate some program terms as part of type checking, and gradual types complicate this in two ways. First, if a gradual language fully embeds an untyped language, then some programs will diverge: indeed, self application ($\lambda x : ?.x\ x$) is typeable in such a language. Second, gradual languages introduce the possibility of type errors that are uncovered as a term is evaluated: applying the function ($\lambda x : ?.x + 1$) may fail, depending on whether its argument can actually be used as a number. So a gradual dependently-typed language must account for the potential of non-termination and failure *during typechecking*.

A gradual dependently-typed language. This work presents GDTL, a gradual dependently-typed core language that supports the whole spectrum between an untyped functional language and a dependently-typed one. As such, GDTL adopts a unified term and type language, meaning that the unknown type $?$ is also a valid term. This allows programmers to specify types with imprecise indices, and to replace proof terms with $?$ (Section 2).

GDTL is a gradual version of the predicative fragment of the Calculus of Constructions with a cumulative universe hierarchy (CC_ω) (Section 3), similar to the core language of Idris [Brady 2013]. We gradualize this language following the Abstracting Gradual Typing (AGT) methodology [Garcia et al. 2016] (Section 4). Because GDTL is a *conservative extension* of this dependently-typed calculus, it is both strongly normalizing and logically consistent for fully static code. These strong properties are however lost as soon as imprecise types and/or terms are introduced. On the dynamic side, GDTL can fully embed the untyped lambda calculus. When writing purely untyped code, static type errors are never encountered. In between, GDTL satisfies the *gradual guarantees* of Siek et al. [2015], meaning that typing and evaluation are monotone with respect to type imprecision. These guarantees ensure that programmers can move code between imprecise and precise types in small, incremental steps, with the program typechecking and behaving identically (modulo dynamic type errors) at each step. If a program fails to typecheck, the programmer knows the problem is not too few type annotations, but rather incompatible types.

GDTL is a call-by-value language with a sharp two-phase distinction. The key technical insight on which GDTL is built is to exploit two distinct notions of evaluation: one for *normalization* during typechecking, and one for *execution* at runtime. Specifically, we present a novel *approximate* normalization technique that guarantees decidable typechecking (Section 5): applying a function of unknown type, which may trigger non-termination, normalizes to the *unknown value*?. Consequently, some terms that would be distinct at runtime become indistinguishable as type indices. Approximation is also used to ensure that compile-time normalization (i.e. during typechecking) always terminates and never signals a “dynamic error”. In this sense, GDTL is closer to Idris than Dependent Haskell, which does admit reduction errors and non-termination during typechecking. At runtime, GDTL uses the standard, precise runtime execution strategy of gradual languages, which may fail due to dynamic type errors, and may diverge as well (Section 6). In that respect, GDTL is closer to Dependent Haskell and Zombie than to Idris, which features a termination checker and a static effect system. We prove that GDTL has decidable typechecking and satisfies all the expected properties of gradual languages [Siek et al. 2015]: type safety, conservative extension of the static language, embedding of the untyped language, and the gradual guarantees (Section 7). We then show how inductive types with eliminators can be added to GDTL without significant changes (Section 8). Section 9 discusses related work, and Section 10 discusses limitations and perspectives for future work.

Disclaimer. This work does not aim to develop a full-fledged dependent gradual type theory, a fascinating objective that would raise many metatheoretic challenges. Rather, it proposes a novel technique, *approximate normalization*, applicable to full-spectrum dependently-typed programming languages. This technique reflects specific design choices that affect the pragmatics of programming and reasoning in GDTL: we review these design decisions, after the informal presentation of the language, in Section 2.5. Also, being a core calculus, GDTL currently lacks several features expected of a practical dependently-typed language (Section 10); nevertheless, this work provides a foundation on which practical gradual dependently-typed languages can be built.

Implementation. We provide a prototype implementation of GDTL in Racket, based on a Redex model. The code for the implementation is open source [Eremondi 2019]. The implementation also supports our extension of natural numbers, vectors, and equality as built-in inductive types.

Technical report. Complete definitions and proofs can be found in [Eremondi et al. 2019] .

2 GOALS AND CHALLENGES

We begin by motivating our goals for GDTL, and describe the challenges and design choices that accompany them.

2.1 The Pain and Promise of Dependent Types

To introduce dependent types, we start with a classic example: length-indexed vectors. In a dependently-typed language, the type $\mathbf{Vec} \ A \ n$ describes any vector that contains n elements of type A . We say that $\mathbf{Vec}$ is *indexed* by the value n . This type has two constructors, $\mathbf{Nil} : (A : \mathbf{Type}_1) \rightarrow \mathbf{Vec} \ A \ 0$ and $\mathbf{Cons} : (A : \mathbf{Type}_1) \rightarrow (n : \mathbf{Nat}) \rightarrow A \rightarrow \mathbf{Vec} \ A \ n \rightarrow \mathbf{Vec} \ A \ (n + 1)$. By making length part of the type, we ensure that operations that are typically partial can only receive values for which they produce results. One can type a function that yields the first element of a vector as follows:

$$\mathbf{head} : (A : \mathbf{Type}_1) \rightarrow (n : \mathbf{Nat}) \rightarrow \mathbf{Vec} \ A \ (n + 1) \rightarrow A$$

Since $\mathbf{head}$ takes a vector of non-zero length, it can never receive an empty vector. The downside of this strong guarantee is that we can only use vectors in contexts where their length is known. This makes it difficult to migrate code from languages with weaker types, or for newcomers to prototype algorithms. For example, a programmer may wish to migrate the following quicksort algorithm into a dependently-typed language:

```
sort vec =      if vec == Nil then Nil else
                (sort (filter (≤ (head vec)) (tail vec)))) ++ head vec
                ++ (sort (filter (> (head vec)) (tail vec)))
```

Migrating this definition to a dependently-typed language poses some difficulties. The recursive calls are not direct deconstructions of $\mathbf{vec}$, so it takes work to convince the type system that the code will terminate, and is thus safe to run at compile time. Moreover, if we try to use this definition with $\mathbf{Vec}$, we must account for how the length of each filtered list is unknown, and while we can prove that the length of the resulting list is the same as the input, this must be done manually. Alternately, we could use simply-typed lists in the dependently-typed language, but we do not wish to duplicate every vector function for lists.

2.2 Gradual Types to the Rescue?

Even at first glance, gradual typing seems like it can provide the desired flexibility. In a gradually-typed language, a programmer can use the unknown type, written $?$, to soften the static typing discipline. Terms with type $?$ can appear in any context. Runtime type checks ensure that dynamically-typed code does not violate invariants expressed with static types.

Since $?$ allows us to embed untyped code in a typed language, we can write a gradually-typed fixed-point combinator $Z : (A : \mathbf{Type}_1) \rightarrow (B : \mathbf{Type}_1) \rightarrow ((A \rightarrow B) \rightarrow A \rightarrow B) \rightarrow A \rightarrow B$. We define this in the same way as the usual Z combinator, but the input function is ascribed the type $?$, allowing for self-application. Using this combinator, the programmer can write $\mathbf{sort}$ using general recursion. Furthermore, the programmer can give $\mathbf{sort}$ the type $? \rightarrow ?$, causing the length of the results of $\mathbf{filter}$ to be ignored in typechecking. Annotating the vector with $?$ inserts runtime checks that ensure that the program will fail (rather than behave in an undefined manner) if it is given an argument that is not a vector.

However, introducing the dynamic type $?$ in a dependently-typed language brings challenges.

The unknown type is not enough. If we assign $\mathbf{sort}$ the type $? \rightarrow ?$, then we can pass it any argument, whether it is a vector or not. This seems like overkill: we want to restrict $\mathbf{sort}$ to vectors and statically rule out nonsensical calls like $\mathbf{sort} \ \mathbf{false}$. Unfortunately the usual notion of type imprecision is too coarse-grained to support this. We want to introduce imprecision more judiciously, as in $\mathbf{Vec} \ A \ ?$: the type of vectors with *unknown length*. But the length is a natural number, not a type. How can we express imprecision in type indices?

Dependent types require proofs. To seamlessly blend untyped and dependently-typed code, we want to let programs omit proof terms, yet still allow code to typecheck and run. But this goes further than imprecision in type indices, since imprecision also manifests in program terms. What should the dynamic semantics of imprecise programs be?

Gradual typing introduces effects. Adding $?$ to types introduces two effects. The ability to type self-applications means programs may diverge, and the ability to write imprecise types introduces the possibility of type errors uncovered while evaluating terms. These effects are troubling because dependent typechecking must often evaluate code, sometimes under binders, to compare dependent types. We must normalize terms at compile time to compute the type of dependent function applications. This means that both effects can manifest *during typechecking*. How should compile-time evaluation errors be handled? Can we make typechecking decidable in the presence of possible non-termination?

Dependent types rely on equality. The key reason for normalizing at compile time is that we must compare types, and since types can be indexed by terms, we need a method of comparing arbitrary terms. If we have $A : (\text{Nat} \rightarrow \text{Nat}) \rightarrow \text{Type}_1$, then $A (\lambda x. 1 + 1)$ and $A (\lambda x. 2)$ should be seen as the same type. In intensional type theories, like that of Coq, Agda and Idris, this is done using *definitional equality*, which fully evaluates terms (even under binders) to determine whether their normal forms are *syntactically* equal. Of course, this is a weaker notion than *propositional equality*, but typechecking with propositional equality (as found in *extensional* theories) is undecidable. In intensional theories, explicit rewriting must be used to exploit propositional equalities.

In a gradual language, types are also compared *at runtime* to compensate for imprecise static type information. With gradual dependent types, how should types (and the terms they contain) be compared at runtime? The simplest solution is to use the same notion of definitional equality as used for static typechecking. This has some unfortunate consequences, such as $A (\lambda x. x + x - x)$ and $A (\lambda x. x)$ being deemed inconsistent, even though they are clearly propositionally equal. However, mirroring compile-time typechecking at runtime simplifies reasoning about the language behavior.

2.3 GDTL in Action

To propagate imprecision to type indices, and soundly allow omission of proof terms, GDTL admits $?$ *both as a type and a term*. To manage effects due to gradual typing, we use separate notions of evaluation for compile-time and runtime. Introducing *imprecision in the compile-time normalization of types* avoids both non-termination and failures during typechecking.

The unknown as a type index. Since full-spectrum dependently-typed languages conflate types and terms, GDTL allows $?$ to be used as either a term or a type. Just as any term can have type $?$, the term $?$ can have any type. This lets dependent type checks be deferred to runtime. For example, we can define vectors `staticNil`, `dynNil` and `dynCons` as follows:

<code>staticNil : Vec Nat 0</code>	<code>dynNil : Vec Nat ?</code>	<code>dynCons : Vec Nat ?</code>
<code>staticNil = Nil Nat</code>	<code>dynNil = Nil Nat</code>	<code>dynCons = Cons Nat 0 0 (Nil Nat)</code>

Then, `(head Nat 1 staticNil)` does not typecheck, `(head Nat 1 dynNil)` typechecks but fails at runtime, and `(head Nat 1 dynCons)` typechecks and succeeds at runtime. The programmer can choose between compile-time or runtime checks, but safety is maintained either way, and in the fully-static case, the unsafe code is still rejected.

The unknown as a term at runtime. Having $?$ as a term means that programmers can use it to optimistically omit *proof terms*. Indeed, terms can be used not only as type indices, but also as proofs of propositions. For example, consider the equality type $\text{Eq} : (A : \text{Type}_i) \rightarrow A \rightarrow A \rightarrow \text{Type}_i$,

along with its lone constructor **Refl** : $(A : \text{Type}_i) \rightarrow (x : A) \rightarrow \text{Eq } A \ x \ x$. We can use these to write a (slightly contrived) formulation of the head function:

$$\text{head}' : (A : \text{Type}_i) \rightarrow (n : \text{Nat}) \rightarrow (m : \text{Nat}) \rightarrow \text{Eq } \text{Nat } n \ (m + 1) \rightarrow \text{Vec } A \ n \rightarrow A$$

This variant accepts vectors of any length, provided the user also supplies a proof that its length n is not zero (by providing the predecessor m and the equality proof $\text{Eq } n \ (m + 1)$). GDTL allows $?$ to be used in place of a proof, while still ensuring that a runtime error is thrown if head' is ever given an empty list. For instance, suppose we define a singleton vector and a proof that $0 = 0$:

$$\begin{aligned} \text{staticCons} &: \text{Vec } \text{Nat } 1 & \text{staticProof} &: \text{Eq } \text{Nat } 0 \ 0 \\ \text{staticCons} &= \text{Cons } \text{Nat } 0 \ (\text{Nil } \text{Nat}) & \text{staticProof} &= \text{Refl } \text{Nat } 0 \end{aligned}$$

Then $(\text{head}' \ \text{Nat } 0 \ 0 \ \text{staticProof } \text{staticNil})$ does not typecheck, $(\text{head}' \ \text{Nat } 0 \ ? \ ? \ \text{staticNil})$ typechecks but fails at runtime, and $(\text{head}' \ \text{Nat } 1 \ ? \ ? \ \text{staticCons})$ typechecks and succeeds at runtime. To see why we get a runtime failure for staticNil , we note that internally, head' uses an explicit *rewriting* of the equality, i.e. if x and y are equal, then any property P that holds for x must also hold for y :

$$\text{rewrite} : (A : \text{Type}_1) \rightarrow (x : A) \rightarrow (y : A) \rightarrow (P : A \rightarrow \text{Type}_1) \rightarrow \text{Eq } A \ x \ y \rightarrow P \ x \rightarrow P \ y$$

In GDTL, when $?$ is treated as an equality proof, it behaves as $\text{Refl } ? \ ?$.¹ Therefore, in the latter two cases of our example, rewrite gives a result of type $P \ ?$, which is checked against type $P \ y$, *at runtime*. If $P \ x$ and $P \ y$ are not definitionally equal, then this check fails with a runtime error.

Managing effects from gradual typing. To illustrate how the effects of gradual typing can show up in typechecking, suppose a programmer uses the aforementioned Z combinator to accidentally write a non-terminating function badFact .

$$\text{badFact} = \lambda m . Z \ (\lambda f . \text{ifzero } m \ (f \ 1)(m * f \ (m))) \quad \text{-- never terminates}$$

As explained before, from a practical point of view, it is desirable for GDTL to fully support dynamically-typed terms, because it allows the programmer to opt out of both the type discipline and the termination discipline of a dependently-typed language. However, this means that computing the return type of a function application may diverge, for instance:

$$\begin{aligned} \text{repeat} &: (A : \text{Type}_1) \rightarrow (n : \text{Nat}) \rightarrow A \rightarrow (\text{Vec } A \ n) \\ \text{factList} &= \text{repeat } \text{Nat} \ (\text{badFact } 1) \ 0 \quad \text{-- has type } \text{Vec } \text{Nat} \ (\text{badFact } 1) \end{aligned}$$

To isolate the non-termination from imprecise code, we observe that any diverging code will necessarily apply a function of type $?$. While badFact does not have type $?$, its definition uses Z , which contains ascriptions of type $?$.

Similarly, a naïve approach to gradual dependent types will encounter failures when normalizing some terms. Returning to our head function, how should we typecheck the following term?

$$\text{faillist} = \text{head } \text{Nat} \ (\text{false} :: ?) \ \text{staticCons}$$

We need to check staticCons against $\text{Vec } \text{Nat} \ ((\text{false} :: ?) + 1)$, but what does $(\text{false} :: ?) + 1$ mean as a vector length?

The difficulty here is that if a term contains type ascriptions that may produce a runtime failure, then it will *always* trigger an error when normalizing, since normalization evaluates under binders.

¹This follows directly from the understanding of the unknown type $?$ as denoting all possible static types [Garcia et al. 2016]. Analogously, the gradual term $?$ denotes all possible static terms. Thus applying the term $?$ as a function represents applying all possible functions, producing all possible results, which can be abstracted as the gradual term $?$. This means that $?$, when applied as a function, behaves as $\lambda x . ?$. Similarly, $?$ treated as an equality proof behaves as $\text{Refl } ? \ ?$.

This means that typechecking will fail any time we apply a function to a possibly-failing term. This is highly undesirable, and goes against the spirit of gradual typing: writing programs in the large would be very difficult if applying a function to an argument that does match its domain type caused a type error, or caused typechecking to diverge! Whereas Dependent Haskell places the burden on the programmer to ensure termination and freedom from failure during typechecking, doing so in a gradual language would make it difficult for programmers because of the possibly indirect interactions with untyped code.

GDTL avoids both problems by using different notions of running programs for the compile-time and runtime phases. We distinguish compile-time *normalization*, which is *approximate but total*, from runtime *execution*, which is *exact but partial*. When non-termination or failures are possible, compile-time normalization uses `?` as an approximate but pure result. So both `factList` and `failList` can be defined and used in runtime code, but they are assigned type `Vec Nat ?`. To avoid non-termination and dynamic failures, we want our language to be strongly normalizing during typechecking. Approximate normalization gives us this.

GDTL normalization is focused around *hereditary substitution* [Watkins et al. 2003], which is a total operation from canonical forms to canonical forms. Because hereditary substitution is structurally decreasing in the *type* of the value being substituted, a static termination proof is easily adapted to GDTL. This allows us to pinpoint exactly where gradual types introduce effects, approximate in those cases, and easily adapt the proof of termination of a static language to the gradual language GDTL. Similarly, our use of bidirectional typing means that a single check needs to be added to prevent failures in normalization.

2.4 Gradual Guarantees for GDTL

To ensure a smooth transition between precise and imprecise typing, GDTL satisfies the *gradual guarantee*, which comes in two parts [Siek et al. 2015]. The *static gradual guarantee* says that reducing the precision of a program preserves its well-typedness. The *dynamic gradual guarantee* states that reducing the precision of a program also preserves its behavior, though the resulting value may be less precise.

One novel insight of GDTL’s design is that the interplay between dependent typechecking and program evaluation carries over to the gradual guarantees. Specifically, the static gradual guarantee fundamentally depends on a restricted variant of the dynamic gradual guarantee. We show that approximate normalization maps terms related by precision to canonical forms related by precision, thereby ensuring that reducing a term’s precision always preserves well-typedness.

By satisfying the gradual typing criteria, and embedding both a fully static and a fully dynamic fragment, GDTL gives programmers freedom to move within the entire spectrum of typedness, from the safety of higher-order logic to the flexibility of dynamic languages. Furthermore, admitting `?` as a term means that we can easily combine code with dependent and non-dependent types, the midpoint between dynamic and dependent types. For example, the simple list type could be written as `List A = Vec A ?`, so lists could be given to vector-expecting code and vice-versa. The programmer knows that as long as vectors are used in vector-expecting code, no crashes can happen, and safety ensures that using a list in a vector operation will always fail gracefully or run successfully. This is significantly different from work on casts to subset types [Tanter and Tabareau 2015] and dependent interoperability [Dagand et al. 2018], where the user must explicitly provide decidable properties or (partial) equivalences.

2.5 Summary of Design Decisions

GDTL embodies several important design decisions, each with tradeoffs related to ease of reasoning and usability of the language.

Fig. 1. SDTL: Syntax and Semantics

By embracing full-spectrum dependent types, GDTL allows types to be first-class citizens: arbitrary terms can appear in types and expressions can produce types as a result. Therefore the programmer does not need to learn a separate index language, and there is no need to recreate term-level operations at the type level.

Sticking to clearly separated phases allows us to adopt different reduction strategies for typechecking and for execution. Crucially, by using *approximate normalization*, we ensure that typechecking in GDTL always terminates: compile-time normalization is a total (though imprecise) operation. This means that some type information is statically lost, with checks deferred to runtime.

GDTL features an unknown term $?$, which resembles term holes in Agda and Idris, and existential variables in Coq; the notable difference is that programs containing $?$ can be run without evaluation getting stuck. Every type in GDTL is therefore inhabited at least by the unknown term $?$, which means that the language is inconsistent as a logic, except for fully-precise programs.

In a gradual language that can embed arbitrary untyped terms, programs may not terminate at runtime. Every type in GDTL contains expressions that can fail or diverge at runtime, due to imprecision. Fully-precise programs are guaranteed to terminate.

Finally, like Coq, Agda, and Idris, GDTL is based on an intensional type theory, meaning that it automatically decides *definitional equality*—i.e. syntactic equality up to normalization—and not propositional equality; explicit rewriting is necessary to exploit propositional equalities. Consequently, runtime checks in GDTL also rely on definitional equality. This makes equality decidable, but means that a runtime error can be triggered even though two (syntactically different) terms are propositionally equal.

3 SDTL: A STATIC DEPENDENTLY-TYPED LANGUAGE

We now present SDTL, a static dependently-typed language which is essentially a bidirectional, call-by-value, cumulative variant of the predicative fragment of CC_ω (i.e. the calculus of constructions with a universe hierarchy [Coquand and Huet 1988]). SDTL is the starting point of our gradualization effort, following the Abstracting Gradual Typing (AGT) methodology [Garcia et al. 2016], refined to accommodate dependent types.

3.1 Syntax and Dynamic Semantics

The syntax of SDTL is shown in Figure 1. Metavariables for the static variants of terms, values, etc. are written in red, sans-serif font. Types and terms share a syntactic category. Functions and applications are in their usual form. Function types are *dependent*: a variable name is given to the argument, and the codomain may refer to this variable. We have a *universe hierarchy*: the lowest types have the type Type_1 , and each Type_i has type Type_{i+1} . This hierarchy is *cumulative*: any value in Type_i is also in Type_{i+1} . Finally, we have a form for explicit type ascriptions.

We use metavariables v, V to range over values, which are the subset of terms consisting only of functions, function types and universes. For evaluation, we use a call-by-value reduction semantics (Figure 1). Ascriptions are dropped when evaluating, and function applications result in (syntactic) substitution. We refer to the values and semantics as *simple* rather than *static*, since they apply equally well to an untyped calculus, albeit without the same soundness guarantees.

3.2 Comparing Types: Canonical Forms

Since dependent types can contain expressions, it is possible that types may contain redexes. Most dependent type systems have a *conversion* rule that assigns an expression type T_1 if it has type T_2 , and T_2 is convertible to T_1 through some sequence of β -conversions, η -conversions, and α -renamings. Instead, we treat types as $\alpha\beta\eta$ -equivalence classes. To compare equivalence classes, we represent them using *canonical forms* [Watkins et al. 2003], denoted with metavariables u and U . These are β -reduced, η -long canonical members of an equivalence class. We compare terms for $\alpha\beta\eta$ -equivalence by normalizing and syntactically comparing their canonical forms.

The syntax for canonical forms is given in Figure 2. We omit well-formedness rules for terms and environments, since the only difference from the typing rules is the η -longness check.

By representing function applications in *spine form* [Cervesato and Pfenning 2003], we can ensure that all heads are variables, and thus no redexes are present, even under binders. The well-formedness of canonical terms is ensured using bidirectional typing [Pierce and Turner 2000]. An *atomic form* can be a universe Type_i , or a variable x applied to 0 or more arguments, which we refer to as its *spine*. Our well-formedness rules ensure the types of atomic forms are themselves atomic. This ensures that canonical forms are η -long, since they cannot have type $(y : U_1) \rightarrow U_2$.

3.3 Typechecking and Normalization

Using the concept of canonical forms, we can now express the type rules for SDTL in Figure 2. To ensure syntax-directedness, we again use bidirectional typing.

The *type synthesis* judgement $\Gamma \vdash t \Rightarrow U$ says that t has type U under context Γ , where the type is treated as an output of the judgement. That is, from examining the term, we can determine its type. Conversely, the *checking* judgement $\Gamma \vdash t \Leftarrow U$ says that, given a type U , we can confirm that t has that type. These rules allow us to propagate the information from ascriptions inwards, so that only top-level terms and redexes need ascriptions.

Most rules in the system are standard. To support dependent types, **SSYNTHAPP** computes the result of applying a particular value. We switch between checking and synthesis using **SSYNTHANN** and **SCHECKSYNTH**. The predicativity of our system is distilled in the **SSYNTHTYPE** rule: Type_i always has type Type_{i+1} . The rule **SCHECKLEVEL** encodes *cumulativity*: we can always treat types as if they were at a higher level, though the converse does not hold. This allows us to check function types against any Type_i in **SCHECKPI**, provided the domain and codomain check against that Type_i .

We distinguish *hereditary substitution* on canonical forms $[u_1/x]^U u_2 = u_3$, from *syntactic substitution* $[x \Rightarrow t_1]t_2 = t_3$ on terms. Notably, the former takes the type of its variable as input, and has canonical forms as both inputs and as output. In **SSYNTHAPP** and **SCHECKPI**, we use the

Fig. 2. SDSL: Canonical Forms and Typing Rules

normalization judgement $\Gamma \vdash u \Leftarrow t \Leftarrow U$, which computes the canonical form of t while checking it against U . Similarly, **SSYNTHANN** uses the judgement $\Gamma \vdash U \Leftarrow T : \text{Type}$, which uses hereditary substitution to compute the canonical form of T while ensuring it checks against some Type_i .

The rules for normalization (Figure 3) directly mirror those for well-typed terms, building up the canonical forms from sub-derivations. In particular, the rule **SNORMSYNTHVAR** η -expands any variables with function types, which allows us to assume that the function in an application will always normalize to a λ -term. (The rules for the eta expansion function $x \rightsquigarrow_\eta u : U$ are standard, so we omit them). We utilize this assumption in **SNORMSYNTHAPP**, where the canonical form of an application is computed using hereditary substitution.

3.4 Hereditary Substitution

Hereditary substitution is defined in Figure 3. At first glance, many of the rules look like a traditional substitution definition. They traverse the expression looking for variables, and replace them with the corresponding term.

However, there are some key differences. Hereditary substitution has canonical forms as both inputs and outputs. The key work takes place in the rule **SHSUBSPINE**. When replacing x with u_1 in $x\bar{s} u_2$, find the substituted forms of u_2 and $x\bar{s}$, which we call u_3 and $\lambda y. u'_1$ respectively. If the inputs are well-typed and η -long, the substitution of the spine will always return a λ -term, meaning that its application to u_3 is not a canonical form. To produce a canonical form in such a case, we continue substituting, recursively replacing y with u_3 in u'_1 . A similar substitution in the codomain of U gives our result type. Thus, if this process terminates, it will always produce a canonical form.

To ensure that the process does, in fact, terminate for well-typed inputs, we define hereditary substitution in terms of the type of the variable being replaced. Since we are replacing a different

$\Gamma \vdash U \Leftarrow T : \text{Type}$

$\Gamma \vdash t \rightsquigarrow u \Rightarrow U \quad | \quad \Gamma \vdash u \Leftarrow y \Leftarrow U$

(Type Normalization with Unknown Level (rules omitted))

(Static Normalization)

SNORMSYNTHANN

$$\frac{\Gamma \vdash U \Leftarrow T : \text{Type} \quad \Gamma \vdash u \Leftarrow t \Leftarrow U}{\Gamma \vdash (t :: T) \rightsquigarrow u \Rightarrow U}$$

SNORMSYNTHVAR

$$\frac{\vdash \Gamma \quad (x : U) \in \Gamma \quad x \rightsquigarrow_{\eta} u : U}{\Gamma \vdash x \rightsquigarrow u \Rightarrow U}$$

SNORMSYNTHAPP

$$\frac{\Gamma \vdash t_1 \rightsquigarrow (\lambda x. u_1) \Rightarrow (x : U_1) \rightarrow U_2 \quad \Gamma \vdash u_2 \Leftarrow t_2 \Leftarrow U_1 \quad [u_2/x]^{U_1} u_1 = u_3 \quad [u_2/x]^{U_1} U_2 = U_3}{\Gamma \vdash t_1 t_2 \rightsquigarrow u_3 \Rightarrow U_3}$$

$[u_1/x]^U u_2 = u_3$

SHSUBPI

$$\frac{[u/x]^U U_1 = U'_1 \quad [u/x]^U U_2 = U'_2 \quad x \neq y}{[u/x]^U (y : U_1) \rightarrow U_2 = (y : U'_1) \rightarrow U'_2}$$

SHSUBDIFFNIL

$$\frac{x \neq y}{[u/x]^U y = y}$$

SHSUBDIFFCONS

$$\frac{x \neq y \quad [u/x]^U y \bar{s} = y \bar{s}' \quad [u/x]^U u_2 = u_3}{[u/x]^U y \bar{s} u_2 = y \bar{s}' u_3}$$

SHSUBTYPE

$$\frac{}{[u/x]^U \text{Type}_i = \text{Type}_i}$$

SHSUBLAM

$$\frac{[u/x]^U u_2 = u_3 \quad x \neq y}{[u/x]^U (\lambda y. u_2) = (\lambda y. u_3)}$$

SHSUBSPINE

$$\frac{[u_1/x]^U x \bar{s} u_2 = u_3 : U'}{[u_1/x]^U x \bar{s} u_2 = u_3}$$

$[u/x]^U x \bar{s} = u' : U'$

(Static Atomic Hereditary Substitution)

SHSUBRSPINE

$$\frac{[u_1/x]^U x \bar{s} = (\lambda y. u'_1) : (y : U'_1) \rightarrow U'_2 \quad [u_1/x]^U u_2 = u_3 \quad [u_3/y]^{U'_1} u'_1 = u'_2 \quad [u_3/y]^{U'_1} U'_2 = U'_3}{[u_1/x]^U x \bar{s} u_2 = u'_2 : U'_3}$$

SHSUBRHEAD

$$\frac{}{[u/x]^U x = u : U}$$

Fig. 3. SDTL: Normalization (select rules) and Hereditary Substitution

variable in the premise **SHSUBRSPINE**, we must keep track of the type of the resultant expression when substituting in spines, which is why substitution on atomic forms is a separate relation. We order types by the multiset of universes of all arrow types that are subterms of the type, similar to techniques used for Predicative System F [Eades and Stump 2010; Mangin and Sozeau 2015]. We can use the well-founded multiset ordering given by Dershowitz and Manna [1979]: if a type U has maximum arrow type universe i , we say that it is greater than all other types containing fewer arrows at universe i whose maximum is not greater than i . Predicativity ensures that, relative to this ordering, the return type of a function application is always less than the type of the function itself. In all premises but the last two of **SHSUBRSPINE**, we recursively invoke substitution on strict subterms, while keeping the type of the variable the same. In the remaining cases, we perform substitution at a type that is smaller by our multiset order.

3.5 Properties of SDTL

Since SDTL is mostly standard, it enjoys the standard properties of dependently-typed languages. Hereditary substitution can be used to show that the language is strongly normalizing, and thus consistent as a logic. Since the type rules, hereditary substitution, and normalization are syntax

$u, U \in \text{GCANONICAL}$		
Gradual Terms	Gradual Canonical Forms	Gradual Atomic Forms
$t, T ::=$	$u, U ::=$	$r, R ::=$
$\lambda x. t$	$\lambda x. u$	$x\bar{s}$
$t_1 t_2$	r	Type_i
x	$?$	
Type_i	$(x : U_1) \xrightarrow{i} U_2$	
$(x : T_1) \rightarrow T_2$		
$t :: T$		
$?$		
		Gradual Canonical Spines
		$\bar{s} ::=$
		$\cdot$
		$\bar{s} u$

Fig. 4. GDTL: Terms and Canonical Forms

directed and terminating, typechecking is decidable. Finally, because all well-typed terms have canonical forms, SDTL is type safe.

4 GDTL: ABSTRACTING THE STATIC LANGUAGE

We now present GDTL, a gradual counterpart to SDTL derived following the Abstracting Gradual Typing (AGT) methodology [Garcia et al. 2016], extended to the setting of dependent types. The key idea behind AGT is that gradual type systems can be designed by first specifying the *meaning* of gradual types in terms of sets of static types. This meaning is given as a concretization function γ that maps a gradual type to the set of static types that it represents, and an abstraction function α that recovers the *most precise* gradual type that represents a given set of static types. In other words, γ and α form a *Galois connection*.

Once the meaning of gradual types is clear, the typing rules and dynamic semantics for the gradual language can be derived systematically. First, γ and α allow us to lift the type predicates and type functions used in the static type system (such as equality, subtyping, join, etc.) to obtain their gradual counterparts. From these definitions, algorithmic characterizations can then be validated and implemented. Second, the gradual type system is obtained from the static type system by using these lifted type predicates and functions. Finally, the runtime semantics follow by proof reduction of the typing derivation, mirroring the type safety argument at runtime. In particular, typing derivations are augmented with pieces of *evidence* for consistent judgments, whose combination during reduction may be undefined, hence resulting in a runtime type error.

In this work we follow the AGT methodology, specifying γ and α , then describing how the typing rules are lifted to gradual types. In doing so, we uncover several points for which the standard AGT approach lacks the flexibility to accommodate full-spectrum dependent types with $?$ as a term. We describe our extensions to (and deviations from) the AGT methodology, and how they allow us to fully support gradual dependent types.

Throughout this section, we assume that we have gradual versions of hereditary substitution and normalization. We leave the detailed development of these notions to Section 5, as they are non-trivial if one wants to preserve both decidable typechecking and the gradual guarantee (Section 2.4). The dynamic semantics of GDTL are presented in Section 6, and its metatheory in Section 7.

4.1 Terms and Canonical Forms

Syntax. The syntax of GDTL (Figure 4) is a simple extension of SDTL's syntax. We use blue, serif font to for metavariables denoting gradual terms, contexts, etc. In addition to constructs from SDTL, GDTL's syntax includes $?$, the unknown term/type. This represents a type or term that is unknown to the programmer: by annotating a type with $?$ or leaving a term as $?$, they can allow their program

$$\begin{array}{l}
\gamma : \text{GCANONICAL} \rightarrow \mathcal{P}(\text{SCANONICAL}) \setminus \emptyset \\
\\
\begin{array}{lll}
\gamma(?) & =_{\text{SCANONICAL}} & \gamma(\lambda x. \mathbf{u}) & = \{ \lambda x. \mathbf{u}' \mid \mathbf{u}' \in \gamma(\mathbf{u}) \} \\
\gamma(\text{Type}_i) & = \{ \text{Type}_i \} & \gamma(x \bar{s} \mathbf{u}) & = \{ x \bar{s}' \mathbf{u}' \mid x \bar{s}' \in \gamma(x \bar{s}), \mathbf{u}' \in \gamma(\mathbf{u}) \} \\
\gamma(x) & = \{ x \} & \gamma((x : \mathbf{U}_1) \rightarrow \mathbf{U}_2) & = \{ (x : \mathbf{U}'_1) \rightarrow \mathbf{U}'_2 \mid \mathbf{U}'_1 \in \gamma(\mathbf{U}_1), \mathbf{U}'_2 \in \gamma(\mathbf{U}_2) \}
\end{array} \\
\\
\boxed{\mathbf{U} \cong \mathbf{U}'} & & & \text{(Consistency of Gradual Canonical Terms)} \\
\\
\begin{array}{ccc}
\text{CONSISTENTEQ} & \text{CONSISTENTPI} & \text{CONSISTENTLAM} \\
\frac{}{\mathbf{u} \cong \mathbf{u}} & \frac{\mathbf{U}_1 \cong \mathbf{U}'_1 \quad \mathbf{U}_2 \cong \mathbf{U}'_2}{(x : \mathbf{U}_1) \rightarrow \mathbf{U}_2 \cong (x : \mathbf{U}'_1) \rightarrow \mathbf{U}'_2} & \frac{\mathbf{u} \cong \mathbf{u}'}{(\lambda x. \mathbf{u}) \cong (\lambda x. \mathbf{u}')} \\
\\
\text{CONSISTENTAPP} & \text{CONSISTENTDYNL} & \text{CONSISTENTDYNR} \\
\frac{x \bar{s} \cong x \bar{s}' \quad \mathbf{u} \cong \mathbf{u}'}{x \bar{s} \mathbf{u} \cong x \bar{s}' \mathbf{u}'} & \frac{}{? \cong \mathbf{u}} & \frac{}{\mathbf{u} \cong ?}
\end{array}
\end{array}$$

Fig. 5. GDTL: Concretization and Consistency

to typecheck with only partial typing information or an incomplete proof. Similarly, $?$ is added to the syntax of canonical values.

Additionally, arrow-types are annotated with a level i . The type $(x : \mathbf{U}_1) \xrightarrow{i} \mathbf{U}_2$ is well formed at type Type_i , and we have a special top level ω where $(x : \mathbf{U}_1) \xrightarrow{\omega} \mathbf{U}_2$ is well formed at type $?$. These annotations are necessary for ensuring the termination of hereditary substitution, but are inferred during normalization, and are never present in source programs. We often omit these annotations, as they clutter the presentation.

Canonical forms do not contain ascriptions. While statically-typed languages use ascriptions only for guiding typechecking, the potential for dynamic type failure means that ascriptions have computational content in gradual typing. Notably, only variables or neutral applications can synthesize $?$ as a type, though any typed expression can be checked against $?$. This allows us to reason about canonical forms at a given type: while we can layer ascriptions on terms, such as $(\text{true} :: ?) :: ? \rightarrow ?$, the only canonical forms with function types are lambdas and $?$.

4.2 Concretization and Predicates

The main idea of AGT is that gradual types abstract sets of static types, and that each gradual type can be made *concrete* as a set of static types. For our system, we simply extend this to say that gradual terms represent sets of static terms. In a simply typed language, a static type embedded in the gradual language concretizes to the singleton set containing itself. However, for terms, we wish to consider the entire $\alpha\beta\eta$ -equivalence class of the static term. As with typechecking, this process is facilitated by considering only canonical forms. The concretization function $\gamma : \text{GCANONICAL} \rightarrow \mathcal{P}(\text{SCANONICAL})$, defined in Figure 5, recurs over sub-terms, with $?$ mapping to the set of all terms.

Given the concretization, we can *lift* a predicate from the static system to the gradual system. A predicate holds for gradual types if it holds for some types in their concretizations. For equality, this means that $\mathbf{U} \cong \mathbf{U}'$ if and only if $\gamma(\mathbf{U}) \cap \gamma(\mathbf{U}') \neq \emptyset$. We present a syntactic version of this in Figure 5. Concretization also gives us a notion of *precision* on gradual types. We say that $\mathbf{U} \sqsubseteq \mathbf{U}'$ if $\gamma(\mathbf{U}) \subseteq \gamma(\mathbf{U}')$: that is, $\mathbf{U}$ is more precise because there are fewer terms it could plausibly represent. We can similarly define $\mathbf{U} \sqcap \mathbf{U}'$ as the most general term that is as precise as both $\mathbf{U}$ and $\mathbf{U}'$. Note

$$\begin{array}{c}
\alpha : \mathcal{P}(\text{SCANONICAL}) \setminus \emptyset \rightarrow \text{GCANONICAL} \\
\alpha(\{x\bar{s} \ u \mid x\bar{s} \in A, u \in B\}) = x\bar{s}' \ u' \quad \text{where } \alpha(A) = x\bar{s}', \alpha(B) = u' \\
\alpha(\{(x : U_1) \rightarrow U_2 \mid U_1 \in A, U_2 \in B\}) = (x : U'_1) \rightarrow U'_2 \quad \text{where } \alpha(A) = U'_1, \alpha(B) = U'_2 \\
\alpha(\{\lambda x. u \mid u \in A\}) = \lambda x. u' \quad \text{where } \alpha(A) = u' \\
\alpha(\{\text{Type}_i\}) = \text{Type}_i \quad \alpha(\{x\}) = x \quad \alpha(A) = ? \text{ otherwise} \\
\boxed{\text{dom} : \text{GCANONICAL} \rightarrow \text{GCANONICAL}} \quad \boxed{[\square/_]\square \text{cod} \square : \text{GCANONICAL}^3 \rightarrow \text{GCANONICAL}} \\
\boxed{[\square/_]\square \text{body} \square : \text{GCANONICAL}^3 \rightarrow \text{GCANONICAL}} \quad \text{(Partial Functions)} \\
\begin{array}{ccc}
\text{DOMAINPI} & \text{CODSUBPI} & \text{BODYSUBPI} \\
\frac{}{\text{dom}(x : U_1) \rightarrow U_2 = U_1} & \frac{\text{CodSUBPI} \quad [u/x]^{U_1} U_2 = U'_2}{[u/_] \text{cod}(x : U_1) \rightarrow U_2 = U'_2} & \frac{\text{BodySUBPI} \quad [u/x]^U u_2 = u'_2}{[u/_] \text{body}(\lambda x. u_2) = u'_2} \\
\text{DOMAINDYN} & \text{CODSUBDYN} & \text{BODYSUBDYN} \\
\frac{}{\text{dom} ? = ?} & \frac{}{[u/_] \text{cod} ? = ?} & \frac{}{[u/_] \text{body} ? = ?}
\end{array}
\end{array}$$

Fig. 6. GDTL: Abstraction and Lifted Functions

that $U \sqcap U'$ is defined if and only if $U \cong U'$ i.e. if $\gamma(U) \cap \gamma(U') \neq \emptyset$, and like the consistency relation, it can be computed syntactically.

4.3 Functions and Abstraction

When typechecking a function application, we must handle the case where the function has type $?$. Since $?$ is not an arrow type, the static version of the rule would fail in all such cases. Instead, we extract the domain and codomain from the type using partial functions. Statically, $\text{dom}(x : U) \rightarrow U' = U$, and is undefined otherwise. But what should the domain of $?$ be?

AGT gives a recipe for lifting such partial functions. To do so, we need the counterpart to concretization: abstraction. The abstraction function α is defined in Figure 6. It takes a set of static terms, and finds the most precise gradual term that is consistent with the entire set. Now, we are able to take gradual terms to sets of static terms, then back to gradual terms. It is easy to see that $\alpha(\gamma(U)) = U$: they are normal forms describing the same equivalence classes. This lets us define our partial functions in terms of their static counterparts: we concretize the inputs, apply the static function element-wise on all values of the concretization for which the function is defined, then abstract the output to obtain a gradual term as a result.

For example, the domain of a gradual term U is $\alpha(\{\text{dom } U' \mid U' \in \gamma(U)\})$, which can be expressed algorithmically using the rules in Figure 6. We define function-type codomains and lambda-term bodies similarly, though we pair these operations with substitution to avoid creating a “dummy” bound variable name for $?$.

Taken together, α and γ form a Galois connection, which ensures that our derived type system is a conservative extension of the static system.

4.4 Typing Rules

Given concretization and abstraction, AGT gives a recipe for converting a static type system into a gradual one, and we follow it closely. Figure 7 gives the rules for typing. Equalities implied by repeated metavariables have been replaced by consistency checks, such as in **GCHECKSYNTH**.

$\boxed{\Gamma \vdash t \Rightarrow U \quad \quad \Gamma \vdash t \Leftarrow U}$		(Well-Typed Gradual Terms)	
$\frac{\text{GSYNTHANN} \quad \Gamma \vdash U \Leftarrow T : \text{Type}_{\Rightarrow i} \quad \Gamma \vdash t \Leftarrow U}{\Gamma \vdash (t :: T) \Rightarrow U}$		$\frac{\text{GSYNTHTYPE} \quad i > 0}{\Gamma \vdash \text{Type}_i \Rightarrow \text{Type}_{i+1}}$	
$\frac{\text{GSYNTHVAR} \quad (x : U) \in \Gamma \quad \vdash \Gamma}{\Gamma \vdash x \Rightarrow U}$		$\frac{\text{GSYNTHDYN}}{\Gamma \vdash ? \Rightarrow ?}$	
$\frac{\text{GSYNTHAPP} \quad \Gamma \vdash t_1 \Rightarrow U \quad \Gamma \vdash u \Leftarrow t_2 \Leftarrow \text{dom } U \quad [u/_]\text{cod } U = U_2}{\Gamma \vdash t_1 t_2 \Rightarrow U_2}$		$\frac{\text{GCHECKPI} \quad \Gamma \vdash U' \Leftarrow T_1 \Leftarrow U \quad U \cong \text{Type} \quad \vdash (x : U')\Gamma \quad (x : U')\Gamma \vdash T_2 \Leftarrow U}{\Gamma \vdash (x : T_1) \rightarrow T_2 \Leftarrow U}$	
$\frac{\text{GCHECKLAMPI} \quad \vdash (x : U_1)\Gamma \quad (x : U_1)\Gamma \vdash t \Leftarrow U_2}{\Gamma \vdash (\lambda x. t) \Leftarrow (x : U_1) \rightarrow U_2}$		$\frac{\text{GCHECKLAMDYN} \quad \vdash (x : ?)\Gamma \quad (x : ?)\Gamma \vdash t \Leftarrow ?}{\Gamma \vdash (\lambda x. t) \Leftarrow ?}$	
$\frac{\text{GCHECKSYNTH} \quad \Gamma \vdash t \Rightarrow U' \quad U' \cong U}{\Gamma \vdash t \Leftarrow U}$		$\frac{\text{GCHECKLEVEL} \quad \Gamma \vdash T \Rightarrow \text{Type}_i \quad 0 < i < j}{\Gamma \vdash T \Leftarrow \text{Type}_j}$	

Fig. 7. GDTL: Type Checking and Synthesis

Similarly, in **GCHECKPI** we use the judgment $U \cong \text{Type}$ to ensure that the given type is consistent to, rather than equal to, some Type_i . Rules that matched on the form of a synthesized type instead use partial functions, as we can see in **GSYNTHAPP**. We split the checking of functions into **GCHECKLAMPI** and **GCHECKLAMDYN** for clarity, but the rules are equivalent to a single rule using partial functions. In **GSYNTHANN**, the judgment $\Gamma \vdash U \Leftarrow T : \text{Type}_{\Rightarrow i}$ denotes *level synthesis*, where we normalize U while inferring at what universe level it resides.

We note that while γ and α are crucial for deriving the definitions of gradual operations, the operations can be implemented algorithmically as syntactic checks; an implementation does not need to compute γ or α . Also, because $\gamma(x) = \{x\}$ for any variable x , consistency, precision and meet are all well-defined on open terms. Consistency corresponds to the gradual lifting of *definitional equality*: $u_1 \cong u_2$ if and only if there is some $u'_1 \in \gamma(u_1)$ and $u'_2 \in \gamma(u_2)$ where $u'_1 =_{\alpha\beta\eta} u'_2$. This reflects our intensional approach: functions are consistent if their bodies are consistent.

We wish to allow the unknown term $?$ to replace any term in a program. But what should its type be? By the AGT philosophy, $?$ represents all terms, so it should synthesize the abstraction of all inhabited types, which is $?$. We encode this in the rule **GSYNTHDYN**. This means that we can use the unknown term in any context.

As with the static system, we represent types in canonical form, which makes consistency checking easy. Well-formedness rules (omitted) are derived from the static system in the same way as the gradual type rules. Additionally, the gradual type rules rely on the *gradual normalization* judgments, $\Gamma \vdash t \rightsquigarrow u \Rightarrow U$ and $\Gamma \vdash u \Leftarrow t \Leftarrow U$, which we explain in [Section 5.3](#).

4.5 Example: Typechecking head of nil

To illustrate how the GDTL type system works, we explain the typechecking of one example from the introduction. Suppose we have types for natural numbers and vectors, and a derivation for $\Gamma \vdash \text{head} : (A : \text{Type}_1) \rightarrow (n : \text{Nat}) \rightarrow \text{Vec } A \ (n + 1) \rightarrow A$. In [Figure 8](#), we show the (partial) derivation of $\Gamma \vdash \text{head Nat } 0 \ ((\text{Nil Nat}) :: \text{Vec Nat } ?) \Rightarrow \text{Nat}$.

The key detail here is that the compile-time consistency check lets us compare 0 to $?$, and then $?$ to 1 , which allows the example to typecheck. Notice how we only check consistency when we

Fig. 8. Type Derivation for head of nil

switch from checking to synthesis. While this code typechecks, it fails at runtime. We step through its execution in [Section 6.4](#).

5 APPROXIMATE NORMALIZATION

In the previous example, normalization was used to compute the type of `head Nat 0`, replacing `n` with `0` in the type of `head`, normalizing `0 + 1` to `1`. This computation is trivial, but not all are. As we saw in [Section 2.3](#), the type-term overlap in GDTL means that code that is run during typechecking may fail or diverge.

A potential solution would be to disallow imprecisely typed code in type indices. However, this approach breaks the criteria for a gradually-typed language. In particular, it would result in a language that violates the static gradual guarantee ([Section 2.4](#)). The static guarantee implies that if a program does not typecheck, the programmer knows that the problem is not the absence of type precision, but that the types present are fundamentally wrong. Increasing precision in multiple places will never cause a program to typecheck if doing so in one place fails.

In this section, we present two versions of gradual substitution. First, we provide *ideal substitution*, which is well defined on all terms, but for which equality is undecidable. Second, we describe *approximate hereditary substitution*, which regains decidability while preserving the gradual guarantee, by producing compile-time canonical forms that are potentially less precise than their runtime counterparts. Thus, we trade precision for a termination guarantee. From this, we build *approximate normalization*, which uses hereditary substitution to avoid non-termination, and avoids dynamic failures by normalizing certain imprecise terms to `?`.

A key insight of this work is that we need separate notions of *compile-time normalization* and *run-time execution*. That is, we use approximate hereditary substitution *only* in our types. Executing our programs at run-time will not lose information, but it may diverge or fail.

For typechecking, the effect of this substitution is that non-equal terms of the unknown type may be indistinguishable at compile-time. Returning to the example from [Section 2.3](#), the user's faulty factorial-length vector will typecheck, but at type `Vec Nat ?`. Using it will never raise a static error due to its length, but it may raise a runtime error.

5.1 Ideal Substitution

Here, we present a definition of gradual substitution for $\alpha\beta\eta$ -equivalence classes of terms. While comparing equivalence classes is undecidable, we will use ideal substitution as the theoretical foundation, showing that our approximate substitution produces the same results as ideal substitution, save for some loss of precision.

The main difficulty with lifting the definition of hereditary substitution is that the set of terms with a canonical form is only closed under hereditary substitution when we assume a static type discipline. The terms `y y` and `λx. x x` are both syntactically canonical, but if we substitute the

$$[u/x]^U x\bar{s} = u_2 : U_2$$

(Approximate Atomic Hereditary Substitution)

$\frac{\text{GHSUBRHEAD}}{[u/x]^U x = u : U}$	$\frac{\text{GHSUBRDYNSPINE} \quad [u/x]^U x\bar{s} = ? : (y : U_1) \rightarrow U_2 \quad [u_2/y]^{U_1} U_2 = U_3}{[u/x]^U x\bar{s} u_2 = ? : U_3}$
$\frac{\text{GHSUBRLAMSPINE} \quad [u/x]^U x\bar{s} = (\lambda y. u_2) : (y : U_1) \rightarrow U_2 \quad [u/x]^U u_1 = u_3 \quad U_1 < U \quad [u_3/y]^{U_1} u_2 = u_4 \quad [u_3/y]^{U_1} U_2 = U_3 \quad u_4 \rightsquigarrow_{\eta} u_5 : U_3}{[u/x]^U x\bar{s} u_1 = u_5 : U_3}$	
$\frac{\text{GHSUBRLAMSPINEORD} \quad [u/x]^U x\bar{s} = (\lambda y. u_2) : (y : U_1) \rightarrow U_2 \quad U_1 \not< U}{[u/x]^U x\bar{s} u_1 = ? : ?}$	$\frac{\text{GHSUBRDYNTYPE} \quad [u_1/x]^U x\bar{s} = u_2 : ?}{[u_1/x]^U x\bar{s} u_2 = ? : ?}$

Fig. 9. GDTL: Approximate Substitution (select rules)

second in for y , there is no normal form. However, both of these terms can be typed in our gradual system. How can $[(\lambda x. x x)/y]^? y y$ be defined?

If we apply the AGT lifting recipe to hereditary substitution, we get a function that may not have a defined output for all gradually well-typed canonical inputs. Even worse is that determining whether substitution is defined for an input is undecidable. By AGT's formulation, $[u/x]^? u'$ would be $\alpha(\{ [u_1/x]^U u_2 \mid u_1 \in \gamma(u), u_2 \in \gamma(u'), U \in \text{SCANONICAL} \})$. To compute α , we must know which of the concretized results are defined, i.e. find all pairs in $\gamma(u) \times \gamma(u')$ for which there exists some U on which static hereditary substitution is defined. This means determining if there is any finite number of substitutions under which the substitution on a (possibly dynamically-typed) term is defined, which requires solving the Halting Problem.

Recall that we introduced canonical forms in Section 3.2 to uniquely represent $\alpha\beta\eta$ -equivalence classes. While canonical forms are not closed under substitutions, equivalence classes are. Going back to our initial example, what we really want is for $[(\lambda x. x x)/y]^? y y$ to be $((\lambda x. x x)(\lambda x. x x))^{\alpha\beta\eta}$, i.e. the set of all terms $\alpha\beta\eta$ -equivalent to Ω .

Thus we define ideal substitution on $\alpha\beta\eta$ -equivalence classes themselves. For this, we do not need hereditary substitution: if $s \in s^{\alpha\beta\eta}$ and $t \in t^{\alpha\beta\eta}$ are terms with their respective equivalence classes, the substitution $[x \mapsto s^{\alpha\beta\eta}]t^{\alpha\beta\eta}$ is simply the equivalence class of $[x \mapsto s]t$. We now have a total operation from equivalence classes to equivalence classes. These classes may have no canonical representative, but the function is defined regardless. If we extend concretization and abstraction to be defined on equivalence classes, this gives us the definition of ideal substitution:

$$[x \mapsto t_1^{\alpha\beta\eta}]t_2^{\alpha\beta\eta} = \alpha(\{ [x \mapsto t'_1]t'_2 \mid t'_1 \in \gamma(t_1), t'_2 \in \gamma(t_2) \})$$

That is, we find the concretization of the gradual equivalence classes, which are sets of static equivalence classes. We then substitute in each combination of these by taking the substitution of a representative element, and abstract over this set to obtain a single gradual equivalence class.

5.2 Approximate Substitution

With a well-defined but undecidable substitution, we now turn to the problem of how to recover decidable comparison for equivalence classes, without losing the gradual guarantees. We again

turn to (gradual) canonical forms as representatives of $\alpha\beta\eta$ -equivalence classes. What happens when we try to construct a hereditary substitution function syntactically, as in SDTL?

The problem is in adapting **SHSUBSPINE**. Suppose we are substituting u for x in $x\bar{s} u_2$, and the result of substituting in $x\bar{s}$ is $(\lambda y. u') : ?$. Following the AGT approach, we can use the **dom** function to calculate the domain of $?$, which is the type at which we substitute y . But this violates the well-foundedness condition we imposed in the static case! Since the domain of $?$ is $?$, eliminating redexes may infinitely apply substitutions without decreasing the size of the type.

In all other cases, we have no problem, since the term we are substituting into is structurally decreasing. So, while equivalence classes give us our ideal, theoretical definition, hereditary substitution provides us with the exact cases we must approximate in order to preserve decidability. To guarantee termination, we must not perform recursive substitutions in spines with $?$ -typed heads.

There are two apparent options for how to proceed without making recursive calls: we either fail when we try to apply a $?$ -typed function, or we return $?$. The former will preserve termination, but it will not preserve the static gradual guarantee. Reducing the precision of a well-typed program's ascriptions should never yield ill-typed code. If applying a dynamically-typed function caused failure, then changing an ascription to $?$ could cause a previously successful program to crash, violating the guarantee.

Our solution is to produce $?$ when applying a function of type $?$. We highlight the changes to hereditary substitution in [Figure 9](#). **GHSUBRDYNTYPE** accounts for $?$ -typed functions, and **GHSUBRDYNSPINE** accounts for $?$ applied as a function.

We must add one more check to guarantee termination, because $?:?$ could be used to circumvent the universe hierarchy. For instance, we can assign $(x : ?) \rightarrow (x \text{ Type}_{99})$ the type Type_1 , and we can even write a version Girard's Paradox [Coquand 1986; Girard 1972] by using $?$ in place of Type . Because of this, **GHSUBRLAMSPINE** manually checks our decreasing metric.

Concretely, $i < \omega$ for every i , and $U < U'$ when the multiset of annotations on arrow types in U is less than that of U' by the well-founded multiset ordering given by [Dershowitz and Manna \[1979\]](#). In the static case, the type of substitution is always decreasing for this metric. In the presence of $?$, we must check if the order is violated and return $?$ if it is, as seen in the rule **GHSUBRLAMSPINEORD**. Unlike applying a function of type $?$, we believe that this case is unlikely to arise in practice unless programmers are deliberately using $?$ to circumvent the universe hierarchy.

5.3 Approximate Normalization

While approximate hereditary substitution eliminates non-termination, we must still account for dynamic failures. We do so with *approximate normalization* ([Figure 10](#)).

To see the issue, consider that we can type the term $(\lambda A. (0 :: ?) :: A)$ against $(A : \text{Type}_1) \rightarrow A$. However, there are no ascriptions in canonical form, since ascriptions can induce casts, which are a form of computation. The term $(\lambda A. 0)$ certainly does not type against $(A : \text{Type}_1) \rightarrow A$, since 0 will not check against the type variable A . However, if we were to raise a type error, we would never be able to apply a function to the above term. In the context $(A : \text{Type}_1)$, the only canonical term with type A is $?$. For the term to have a well-typed normal form, its body must be $?$.

More broadly, normalization does not preserve *synthesis* of typing, only checking. In the rule **GNCHECKSYNTH**, if $\Gamma \vdash t \rightsquigarrow u \Rightarrow U'$, then the normal form of t might check against U , but it won't necessarily synthesize U (or any type). We need to construct a canonical form u for t at type U , assuming we have some normal form u' for t at type U' . If $U \sqsubseteq U'$, u' will check against U . Otherwise, the only value we can be sure will check against U is $?$, which checks against any type. We formalize this using a pair of rules: **GNCHECKSYNTH** normalizes fully when we can do so in a type-safe way, and **GNCHECKAPPROX** produces $?$ as an approximate result in all other cases.

<div style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> $\Gamma \vdash t \rightsquigarrow u \Rightarrow U \quad \quad \Gamma \vdash u \Leftarrow t \Leftarrow U$ </div> <div style="text-align: center;">GNSYNTHAPP</div> $\frac{\begin{array}{c} \Gamma \vdash t_1 \rightsquigarrow u_1 \Rightarrow U \quad u_1 \rightsquigarrow_{\eta} u'_1 : ? \rightarrow ? \\ \text{dom } U = U_1 \quad \Gamma \vdash u_2 \Leftarrow t_2 \Leftarrow U_1 \\ [u_2/__]^{U_1} \text{body } u'_1 = u_3 \\ [u_2/__] \text{cod } U = U_2 \quad u_3 \rightsquigarrow_{\eta} u'_3 : U_2 \end{array}}{\Gamma \vdash t_1 t_2 \rightsquigarrow u'_3 \Rightarrow U_2}$ <div style="text-align: center;">GNCHECKPiTYPE</div> $\frac{\begin{array}{c} \Gamma \vdash U_1 \Leftarrow T_1 \Leftarrow \text{Type}_i \\ \vdash (x : U_1) \Gamma \\ (x : U_1) \Gamma \vdash U_2 \Leftarrow T_2 \Leftarrow \text{Type}_i \end{array}}{\Gamma \vdash (x : U_1) \xrightarrow{i} U_2 \Leftarrow (x : T_1) \rightarrow T_2 \Leftarrow \text{Type}_i}$	<div style="text-align: right;">(Approximate Normalization)</div> <div style="text-align: center;">GNCHECKSYNTH</div> $\frac{\begin{array}{c} \Gamma \vdash t \rightsquigarrow u \Rightarrow U' \\ U \cong U' \quad U' \sqsubseteq U \end{array}}{\Gamma \vdash u \Leftarrow t \Leftarrow U}$ <div style="text-align: center;">GNCHECKAPPROX</div> $\frac{\begin{array}{c} \Gamma \vdash t \rightsquigarrow u \Rightarrow U' \\ U \cong U' \quad U' \not\sqsubseteq U \end{array}}{\Gamma \vdash ? \Leftarrow t \Leftarrow U}$ <div style="text-align: center;">GNCHECKPiDYN</div> $\frac{\begin{array}{c} \Gamma \vdash U_1 \Leftarrow T_1 \Leftarrow ? \\ \vdash (x : U_1) \Gamma \\ (x : U_1) \Gamma \vdash U_2 \Leftarrow T_2 \Leftarrow ? \end{array}}{\Gamma \vdash (x : U_1) \xrightarrow{\omega} U_2 \Leftarrow (x : T_1) \rightarrow T_2 \Leftarrow ?}$
---	--

Fig. 10. GDTL: Approximate Normalization (select rules)

Gradual typing must also treat η -expansion carefully. We η -expand all variables in **GNSYNTHVAR**, but in **GNSYNTHAPP**, we may be applying a function of type $?$. However, a variable x of type $?$ is η -long. Since we are essentially treating a value of type $?$ as type $? \rightarrow ?$, we must expand x to $(\lambda y. x y)$. We do this in **GNSYNTHVAR** with an extra η -expansion at type $? \rightarrow ?$, which expands a $?$ -typed term one level, and has no effect on a canonical form with a function type.

Normalization is also where we generate the annotations necessary for ensuring the decreasing metric of hereditary substitution. As we see in the rules **GNCHECKPiTYPE** and **GNCHECKPiDYN**, we annotate arrows either with the level against which they are checked, or with ω when checking against $?$. The remaining rules for normalization (omitted) directly mirror the rules from Figure 7. **Type_i**, $?$, and variables all normalize to themselves, and all other rules simply construct normal forms from the normal forms of their subterms.

Some of the difficulty with normalization arises because function arguments are normalized before being substituted. One could imagine a language that normalizes after substituting function arguments, and typechecking fails if a dynamic error is encountered during normalization. Here, normalization could fail, but only on terms that had truly ill-formed types, since unused failing values would be discarded. We leave the development of such a language to future work.

5.4 Properties of Approximate Normalization

Relationship to the Ideal. If we expand our definition of concretization to apply to equivalence classes of terms, gradual precision gives us a formal relationship between ideal and approximate normalization:

THEOREM 5.1 (NORMALIZATION APPROXIMATES THE IDEAL). *For any Γ, t, U , if $\Gamma \vdash t \Leftarrow U$, then $\Gamma \vdash u \Leftarrow t \Leftarrow U$ for some u , and $t^{\alpha\beta\eta} \sqsubseteq u$.*

Intuitively, this holds because approximate normalization for a term either matches the ideal, or produces $?$, which is less precise than every other term.

Preservation of Typing. To prove type safety for GDTL, a key property of normalization is that it preserves typing. This property relies on the fact that hereditary substitution preserves typing, which can be shown using a technique similar to that of Pfenning [2008].

Fig. 11. Evidence Term Syntax and Typing (select rules)

THEOREM 5.2 (NORMALIZATION PRESERVES TYPING). *If $\Gamma \vdash u \leftarrow t \Leftarrow U$, then $\Gamma \vdash u \Leftarrow U$.*

Normalization as a Total Function. Since we have defined substitution and normalization using inference rules, they are technically relations rather than functions. Since the rules are syntax directed in terms of their inputs, it is easy to show that there is at most one result for every set of inputs. As we discussed above, the approximation in **GNCHECKAPPROX** makes normalization total.

THEOREM 5.3 (NORMALIZATION IS TOTAL). *If $\Gamma \vdash t \Leftarrow U$, then $\Gamma \vdash u \leftarrow t \Leftarrow U$ for exactly one u .*

6 GDTL: RUNTIME SEMANTICS

With the type system for GDTL realized, we turn to its dynamic semantics. Following the approaches of **Garcia et al. [2016]** and **Toro et al. [2018b]**, we let the syntactic type-safety proof for the static SDTL drive its design. In place of a cast calculus, gradual terms carry *evidence* that they match their type, and computation steps evolve that evidence incrementally. When evidence no longer supports the well-typedness of a term, execution fails with a runtime type error.

6.1 The Runtime Language

Figure 11 gives the syntax for our runtime language. It mirrors the syntax for gradual terms, with two main changes. In place of type ascriptions, we have a special form for terms augmented with evidence, following **Toro et al. [2018b]**. We also have **err**, an explicit term for runtime type errors.

Translation proceeds by augmenting our bidirectional typing rules to output the translated term. Type ascriptions are dropped in the **GSYNTHANN** rule, and initial evidence of consistency is added in **GCHECKSYNTH**. **Section 6.2** describes how to derive this initial evidence. In the **GSYNTHDYN** rule, we annotate **?** with evidence **?**, so **?** is always accompanied by some evidence of its type. Similarly, functions of type **?** are ascribed $\langle ? \rightarrow ? \rangle$.

In **Figure 11** we also define the class of syntactic values, which determines those terms that are done evaluating. We wish to allow values to be augmented with evidence, but not to have multiple

evidence objects stacked on a value. To express this, we separate the class of values from the class of *raw values*, which are never ascribed with evidence at the top level.

Values are similar to, but not the same, as canonical forms. In particular, there are no redexes in canonical terms, even beneath a λ , whereas values may contain redexes within abstractions.

6.2 Typing and Evidence

To establish progress and preservation, we need typing rules for evidence terms, whose key rules we highlight in Figure 11. These are essentially the same as for gradual terms, with two major changes. First, we no longer use bidirectional typing, since our type system need not be syntax directed to prove safety. Second, whereas gradual terms could be given any type that is consistent with their actual type, we only allow this for dynamic terms directly ascribed with evidence, as seen in the rule **EVTYPEEV**. Thus, all applications of consistency are made explicit in the syntax of evidence terms, and for a term $\varepsilon \mathbf{e}$, the evidence ε serves as a concrete witness between the actual type of $\mathbf{e}$ and the type at which it is used. The normalization relation is extended to evidence terms: it simply erases evidence ascriptions and otherwise behaves like the normalization for gradual terms. We can then define hereditary substitutions of evidence terms into types, which is crucial for updating evidence after function applications.

This raises the question: what is evidence? At a high level, the evidence attached to a term tracks the most precise type information about this term that is dynamically available. As we can see in Figure 11, evidence consists of a canonical type: we use brackets $\langle \rangle$ to syntactically distinguish evidence from canonical forms. Ascribing a term with evidence behaves like a cast in a gradual cast calculus; the key difference is that evidence only ever increases in precision. It serves as a witness for the consistency of two types, and by refining evidence at each step (and failing when this is not possible), we ensure that each intermediate expression is (gradually) well-typed. Garcia et al. [2016] identify the correspondence between evidence and the middle type in a threesome calculus [Siek and Wadler 2010].

AGT provides a general formulation of evidence, applicable to multi-argument, asymmetric predicates. However, since equality is the only predicate we use, the meet of two terms is sufficient to serve as evidence of their consistency. We say that $\varepsilon \vdash \mathbf{U}_1 \cong \mathbf{U}_2$ whenever $\varepsilon = \langle \mathbf{U}' \rangle$ and $\mathbf{U}' \sqsubseteq \mathbf{U}_1 \sqcap \mathbf{U}_2$. There are two key operations on evidence. First, we need *initial evidence* for elaborating gradual terms to evidence terms. If a term synthesizes some $\mathbf{U}$ and is checked against $\mathbf{U}'$, then during elaboration we can ascribe to it the evidence $\langle \mathbf{U} \sqcap \mathbf{U}' \rangle$. Secondly, we need a way to combine two pieces of evidence at runtime, an operation referred to as *consistent transitivity* in AGT: if $\langle \mathbf{U} \rangle \vdash \mathbf{U}_1 \cong \mathbf{U}_2$, and $\langle \mathbf{U}' \rangle \vdash \mathbf{U}_2 \cong \mathbf{U}_3$, then $\langle \mathbf{U} \sqcap \mathbf{U}' \rangle \vdash \mathbf{U}_1 \cong \mathbf{U}_3$, provided that the meet is defined. So we can also use the precision meet to dynamically combine different pieces of evidence.

Evidence is combined using the meet operation, which is based on definitional (intensional) equality. This means that if we have a type $A : (\text{Nat} \rightarrow \text{Nat}) \rightarrow \text{Type}_1$, then $A(\lambda x. x + x - x)$ and $A(\lambda x. x)$ will not be consistent at runtime, despite being extensionally equivalent. Extensional equality is undecidable, so it cannot be used during typechecking. Since definitional equality is decidable, we use it both for typechecking and at runtime. This also ensures that the type operations performed at runtime directly mirror those performed during typechecking.

6.3 Developing a Safe Semantics

To devise our semantics, we imagine a hypothetical proof of progress and preservation. Progress tells us which expressions we need reduction rules for, and preservation tells us how to step in order to remain well-typed.

$$\boxed{e_1 \longrightarrow e_2}$$

(Evidence-based Small-Step Semantics)

$$\frac{\text{STEPASCR} \quad \varepsilon_1 \sqcap \varepsilon_2 = \varepsilon_3}{\varepsilon_1 (\varepsilon_2 \mathbf{w}) \longrightarrow \varepsilon_3 \mathbf{w}}$$

$$\frac{\text{STEPASCRFAIL} \quad \varepsilon_1 \sqcap \varepsilon_2 \text{ undefined}}{\varepsilon_1 (\varepsilon_2 \mathbf{w}) \longrightarrow \text{err}}$$

$$\frac{\text{STEPAPPDYN} \quad \cdot \vdash \mathbf{u} \leftarrow \mathbf{v} \Leftarrow \text{dom } \mathbf{U} \quad [\mathbf{u}/_]\text{cod } \mathbf{U} = \mathbf{U}_2}{(\langle \mathbf{U} \rangle ?) \mathbf{v} \longrightarrow \langle \mathbf{U}_2 \rangle ?}$$

$$\frac{\text{STEPAPPEV} \quad \mathbf{U}' \sqcap \text{dom } \mathbf{U} = \mathbf{U}_1 \quad \cdot \vdash \mathbf{u} \leftarrow \mathbf{w} \Leftarrow \mathbf{U}_1 \quad [\mathbf{u}/_]\text{cod } \mathbf{U} = \mathbf{U}_2}{(\langle \mathbf{U} \rangle (\lambda x. \mathbf{e})) (\langle \mathbf{U}' \rangle \mathbf{w}) \longrightarrow \langle \mathbf{U}_2 \rangle ([x \mapsto \langle \mathbf{U}_1 \rangle \mathbf{w}]^{\mathbf{u}:\mathbf{U}_1} \mathbf{e})}$$

$$\frac{\text{STEPAPPEVRAW} \quad \cdot \vdash \mathbf{u} \leftarrow \mathbf{w} \Leftarrow \text{dom } \mathbf{U} \quad [\mathbf{u}/_]\text{cod } \mathbf{U} = \mathbf{U}_2}{(\langle \mathbf{U} \rangle (\lambda x. \mathbf{e})) \mathbf{w} \longrightarrow \langle \mathbf{U}_2 \rangle ([x \mapsto (\langle \mathbf{U}_1 \rangle \mathbf{w})]^{\mathbf{u}:\mathbf{U}_1} \mathbf{e})}$$

$$\frac{\text{STEPAPPFAILTRANS} \quad \text{dom } \mathbf{U} \sqcap \mathbf{U}' \text{ undefined}}{(\langle \mathbf{U} \rangle (\lambda x. \mathbf{e})) (\langle \mathbf{U}' \rangle \mathbf{w}) \longrightarrow \text{err}}$$

$$\frac{\text{STEPCONTEXT} \quad e_1 \longrightarrow e_2 \quad e_1, e_2 \neq \text{err}}{C[e_1] \longrightarrow C[e_2]}$$

$$\frac{\text{STEPCONTEXTERR} \quad e \longrightarrow \text{err}}{C[e] \longrightarrow \text{err}}$$

Fig. 12. GDTL: Dynamic Semantics

Double Evidence. Since values do not contain terms of the form $\varepsilon_2 (\varepsilon_1 \mathbf{w})$, progress dictates that we need a reduction rule for such a case. If $\cdot \vdash \mathbf{w} : \mathbf{U}$, $\varepsilon_1 \vdash \mathbf{U} \cong \mathbf{U}'$ and $\varepsilon_2 \vdash \mathbf{U}' \cong \mathbf{U}''$, then $\varepsilon_1 \sqcap \varepsilon_2 \vdash \mathbf{U} \cong \mathbf{U}''$, so we step to $(\varepsilon_1 \sqcap \varepsilon_2) \mathbf{w}$. If the meet is not defined, then a runtime error occurs.

Functions with Evidence. There are two complications for reducing applications with evidence. The first is that in $\lambda x. \mathbf{e}$, the variable x may be free in evidence ascriptions within $\mathbf{e}$. When performing a substitution, we need the type and normal form of the term replacing the variable. We use the notation $[x \mapsto \mathbf{e}_1]^{\mathbf{u}:\mathbf{U}} \mathbf{e}_2$ to denote the syntactic replacement of x by $\mathbf{e}_1$ in $\mathbf{e}_2$, where free occurrences of x in evidence within $\mathbf{e}_2$ are replaced by $\mathbf{u}$ (the normal form of $\mathbf{e}_2$) using hereditary substitution at type $\mathbf{U}$. We use this operation to reduce applications.

A second issue is that, while the simple rules dictate how to evaluate a λ -term applied to a value, they do not determine how to proceed for applications of the form $(\langle \mathbf{U} \rangle \lambda x. \mathbf{e}) (\langle \mathbf{U}' \rangle \mathbf{w})$. In such a case, we know that $\cdot \vdash (\langle \mathbf{U} \rangle \lambda x. \mathbf{e}) : \mathbf{U}_1$ and that $\langle \mathbf{U} \rangle \vdash \mathbf{U}_1 \cong \mathbf{U}_2$ for some $\mathbf{U}_2$. Computing $(\text{dom } \langle \mathbf{U} \rangle) \sqcap \langle \mathbf{U}' \rangle$ yields evidence that the type of $\mathbf{w}$ is consistent with the domain of $\mathbf{U}_1$, so we ascribe this evidence during substitution to preserve well-typedness. The evidence-typing rules say that the type of an application is found by normalizing the argument value and substituting into the codomain of the function type. To produce a result at this type, we can normalize $\mathbf{w}$ and substitute it into the codomain of $\langle \mathbf{U} \rangle$, thereby producing evidence that the actual result is consistent with the return type. In the case where $\mathbf{w}$ is not ascribed with evidence, we can simply behave as if it were ascribed $\langle ? \rangle$ and proceed using the above process.

Applying The Unknown Term. The syntax for values only admits application under binders, so we must somehow reduce terms of the form $(\varepsilon ?) \mathbf{v}$. The solution is simple: if the function is unknown, so is its output. Since the unknown term is always accompanied by evidence at runtime, we calculate the result type by substituting the argument into the codomain of the evidence associated with $?$.

The Full Semantics. All other well-typed terms are either values, or contain a redex as a subterm, either of the simple variety or of the varieties described above. Using contextual rules to account for

Proc. ACM Program. Lang., Vol. 3, No. ICFP, Article 88. Publication date: August 2019.

these remaining cases, we have a semantics that satisfies progress and preservation *by construction*. Figure 12 gives the full set of rules.

6.4 Example: Running head of nil

We return to the example from Section 4.5, this time explaining its runtime behaviour. Because of consistency, the term $(\text{Nil Nat} :: (\text{Vec Nat } ?))$ is given the evidence $\langle \text{Vec Nat } ? \rangle$, obtained by computing $\langle \text{Vec Nat } 0 \sqcap \text{Vec Nat } ? \rangle$. Applying consistency to use this as an argument adds the evidence $\langle \text{Vec Nat } 1 \rangle$, since we check $(\text{Nil Nat} :: (\text{Vec Nat } ?))$ against $\text{dom } (\text{Vec Nat } 1 \rightarrow \text{Nat})$. The rule **STEPCONTEXT** dictates that we must evaluate the argument to a function before evaluating the application itself. Our argument is $\langle \text{Vec Nat } 1 \rangle \langle \text{Vec Nat } 0 \rangle (\text{Nil Nat})$, and since the meet of the evidence types is undefined, we step to err with **STEPASCRFAIL**.

7 PROPERTIES OF GDTL

GDTL satisfies all the criteria for gradual languages set forth by Siek et al. [2015].

Safety. First, GDTL is type safe *by construction*: the runtime semantics are specifically crafted to maintain progress and preservation. We can then obtain the standard safety result for gradual languages, namely that well-typed terms do not get stuck.

THEOREM 7.1 (TYPE SAFETY). *If $\vdash e : U$, then either $e \rightarrow^* v$ for some v , $e \rightarrow^* \text{err}$, or e diverges.*

This means that gradually well-typed programs in GDTL may fail with runtime type errors, but they will never get stuck. Among the three main approaches to deal with gradual types in the literature, GDTL follows the original approach of Siek and Taha [2006] and Siek et al. [2015], which enforces types eagerly at boundaries, including at higher-order types. This is in contrast with first-order enforcement (a.k.a as transient semantics [Vitousek et al. 2017]), or simple type erasure (a.k.a as optional typing).² In particular, while the transient semantics support open world soundness [Vitousek et al. 2017] when implemented on top of a (safe) dynamic language, it is unclear if and how this approach, which is restricted to checking type constructors, can scale to full-spectrum dependent types. GDTL is a sound gradually-typed language that requires elaboration of the complete program in order to insert the pieces of evidence that support runtime checking.

Conservative Extension of SDTL. It is easy to show that GDTL is a conservative extension of SDTL. This means that any fully-precise GDTL programs enjoy the soundness and logical consistency properties that SDTL guarantees. Any statically-typed term is well-typed in GDTL by construction, thanks to AGT: on fully precise gradual types, $\alpha \circ \gamma$ is the identity. Moreover, the *only* additions are those pertaining to $?$, meaning that if we restrict ourselves to the static subset of terms (and types) without $?$, then we have all the properties of the static system. We formalize this as follows:

THEOREM 7.2. *If Γ, t, U are the embeddings of some Γ, t, U into GDTL, and $\Gamma \vdash t \Leftarrow U$, then $\Gamma \vdash t \Leftarrow U$. Moreover, if $t \rightarrow^* v$ and t elaborates to e , then there exists some v where $e \rightarrow^* v$, where removing evidence from v yields v .*

Embedding of Untyped Lambda Calculus. A significant property of GDTL is that it can fully embed the untyped lambda calculus, including non-terminating terms. Given an untyped embedding function $[t]$ that (in essence) annotates all terms with $?$ we can show that any untyped term can be embedded in our system. Since no type information is present, all evidence objects are formed using $?$ or $\rightarrow$, and the meet operator never fails and untyped programs behave normally in GDTL.

²Greenman and Felleisen [2018] present a detailed comparative semantic account of these three approaches.

THEOREM 7.3. *For any untyped λ -term t and closing environment Γ that maps all variables to type $?$, then $\Gamma \vdash [t] \Rightarrow ?$. Moreover, if t is closed, then $t \longrightarrow^* v$ implies that $[t]$ elaborates to e where $e \longrightarrow^* v$ and stripping evidence from v yields v .*

Gradual Guarantees. GDTL smoothly supports the full spectrum between dependent and untyped programming—a property known as the gradual guarantee [Siek et al. 2015], which comes in two parts. We say that $\Gamma \sqsubseteq \Gamma'$ if they contain the same variables, and for each $(x : U) \in \Gamma$, $(x : U') \in \Gamma'$ where $U \sqsubseteq U'$.

THEOREM 7.4 (GRADUAL GUARANTEE).

(*STATIC GUARANTEE*) Suppose $\Gamma \vdash t \Leftarrow U$ and $U \sqsubseteq U'$. If $\Gamma \sqsubseteq \Gamma'$ and $t \sqsubseteq t'$, then $\Gamma' \vdash t' \Leftarrow U'$.

(*DYNAMIC GUARANTEE*) Suppose that $\cdot \vdash e_1 : U$, $\cdot \vdash e'_1 : U'$, $e_1 \sqsubseteq e'_1$, and $U \sqsubseteq U'$. If $e_1 \longrightarrow^* e_2$, then $e'_1 \longrightarrow^* e'_2$ where $e_2 \sqsubseteq e'_2$.

AGT ensures that the gradual guarantee holds by construction. Specifically, because approximate normalization and consistent transitivity are monotone with respect to precision, we can establish a weak bisimulation between the steps of the more and less precise versions [Garcia et al. 2016].

A novel insight that arises from our work is that we need a restricted form of the dynamic gradual guarantee *for normalization* in order to prove the static gradual guarantee. To differentiate it from the standard one, we call it the *normalization gradual guarantee*. Because an η -long term might be longer at a more precise type, we phrase the guarantee modulo η -equivalence: we say that $U_1 \sqsubseteq^\eta U_2$ if $U_1 =_\eta U'_1$, $U_2 =_\eta U'_2$ and $U'_1 \sqsubseteq U'_2$.

With these defined, we can state the normalization gradual guarantee:

LEMMA 7.5 (NORMALIZATION GRADUAL GUARANTEE). Suppose $\Gamma_1 \vdash u_1 \Leftarrow t_1 \Leftarrow U_1$. If $\Gamma_1 \sqsubseteq^\eta \Gamma_2$, $t_1 \sqsubseteq t_2$, and $U_1 \sqsubseteq^\eta U_2$, then $\Gamma_2 \vdash u_2 \Leftarrow t_2 \Leftarrow U_2$ where $u_1 \sqsubseteq^\eta u_2$.

8 EXTENSION: INDUCTIVE TYPES

Though GDTL provides type safety and the gradual guarantees, its lack of inductive types means that programming is cumbersome. Church encodings allow for some induction, but are strictly less powerful than proper inductive types. Additionally, induction principles, along with basic facts like $0 \neq 1$, cannot be proven in the purely negative Calculus of Constructions [Stump 2017]. However, we can type such a term if we introduce inductive types with eliminators, and allow types to be defined in terms of such eliminations.

This section describes how to extend GDTL with a few common inductive types—natural numbers, vectors, and an identity type for equality proofs—along with their eliminators. While not as useful as user-defined types or pattern matching (both of which are important subjects for future work), this specific development illustrates how our approach can be extended to a more full-fledged dependently-typed language. Note that while we show how inductives can be added to the language, extending our metatheory to include inductives is left as future work.

Syntax and Typing. We augment the syntax for terms as follows:

$$t, T ::= \dots \mid \text{Nat} \mid 0 \mid \text{Succ } t \mid \text{Vec } T \ t \mid \text{Eq } T \ t_1 \ t_2 \mid \text{Refl } T \ t \mid \text{Nil } t \mid \text{Cons } T \ t_1 \ t_2 \ t_3 \\ \mid \text{natElim } T \ t_1 \ t_2 \ t_3 \mid \text{vecElim } T_1 \ t_1 \ T_2 \ t_2 \ t_3 \ t_4 \mid \text{eqElim } T_1 \ T_2 \ t_1 \ t_2 \ t_3 \ t_4$$

The typing rules are generally straightforward. We omit the full rules, but we essentially type them as functions that must be fully applied, with the types given in Figure 13. Each form checks its arguments against the specified types, and the rule **GCHECKSYNTH** ensures that typechecking succeeds so long as argument types are consistent with the expected types. Adding these constructs to canonical forms is interesting. Specifically, the introduction forms are added as atomic forms,

$$\begin{array}{ll}
\text{Nat} : \text{Type}_1 & 0 : \text{Nat} \\
\text{Vec} : \text{Type}_i \rightarrow \text{Nat} \rightarrow \text{Type}_i & \text{Succ} : \text{Nat} \rightarrow \text{Nat} \\
\text{Eq} : (A : \text{Type}_i) \rightarrow A \rightarrow A \rightarrow \text{Type}_i & \text{Nil} : (A : \text{Type}_i) \rightarrow \text{Vec } A \ 0 \\
& \text{Refl} : (A : \text{Type}_i) \rightarrow (x : A) \rightarrow \text{Eq } A \ x \ x \\
\text{Cons} : (A : \text{Type}_i) \rightarrow (n : \text{Nat}) \rightarrow (\text{hd} : A) \rightarrow (\text{tl} : \text{Vec } A \ n) \rightarrow \text{Vec } A \ (\text{Succ } n) \\
\text{natElim} : (m : (\text{Nat} \rightarrow \text{Type}_i)) \rightarrow (m \ 0) \rightarrow ((n : \text{Nat}) \rightarrow m \ n \rightarrow m \ (\text{Succ } n)) \rightarrow (n : \text{Nat}) \rightarrow m \ n \\
\text{vecElim} : (A : \text{Type}_i) \rightarrow (n : \text{Nat}) \rightarrow (m : (\text{Vec } A \ n \rightarrow \text{Type}_i)) \rightarrow m \ (\text{Vec } A \ 0) \\
\quad \rightarrow ((n_2 : \text{Nat}) \rightarrow (h : A) \rightarrow (\text{tl} : \text{Vec } A \ n_2) \rightarrow m \ \text{vec} \rightarrow m \ (\text{Cons } A \ (\text{Succ } n_2) \ \text{hd} \ \text{tl})) \\
\quad \rightarrow (\text{vec} : \text{Vec } A \ n) \rightarrow m \ \text{vec} \\
\text{eqElim} : (A : \text{Type}_i) \rightarrow (m : (x : A) \rightarrow (y : A) \rightarrow \text{Eq } A \ x \ y \rightarrow \text{Type}_i) \\
\quad \rightarrow ((z : A) \rightarrow m \ z \ z \ (\text{Refl } A \ z)) \rightarrow (x : A) \rightarrow (y : A) \rightarrow (p : \text{Eq } A \ x \ y) \rightarrow (m \ x \ y \ p)
\end{array}$$

Fig. 13. Constructor and Eliminator Types

and the eliminators become new variants of the canonical spines. Since **natElim** applied to a **Nat** is a redex, canonical forms can eliminate variables. Eliminators take one fewer argument than in the term version, since the last argument is always the rest of the spine in which the eliminator occurs.

$$\begin{array}{ll}
r, R ::= & \dots \mid \text{Nat} \mid 0 \mid \text{Succ } u \mid \text{etc} \dots \\
\bar{s} ::= & \dots \mid \text{natElim } u_1 \ u_2 \ u_3 \mid \text{vecElim } U_1 \ u_1 \ U_2 \ u_2 \ u_3 \mid \text{etc} \dots
\end{array}$$

Normalization. We extend hereditary substitution to inductive types. Unfortunately, we must treat hereditary substitution as a *relation* between normal forms. The strictly-decreasing metric we previously used no longer holds for inductive types, so we have not proved that hereditary substitution with inductives is a well-defined function; this is left as future work. For introduction forms, we simply substitute in the subterms. For eliminators, if we are ever replacing x with u' in $x\bar{s} \text{ natElim } u_1 \ u_2 \ u_3$, then we substitute in $x\bar{s}$ and see if we get 0 or **Succ** back. If we get 0, we produce u_2 , and if we get **Succ** n , we compute the recursive elimination for n as u'_2 , and substitute n and u'_2 for the arguments of u_3 . Vectors are handled similarly. An application of **eqElim** to **Refl** $A \ x$ simply returns the given value of type $m \ x \ x \ (\text{Refl } x \ A)$ as a value of type $m \ x \ y \ p$.

How should we treat eliminations with $?$ as a head? Since $?$ represents the set of all static values, the result of eliminating it is the abstraction of the eliminations of all possible values. Since these values may produce conflicting results, the abstraction is simply $?$, which is our result. However, for equality, we have a special case. Each instance of **eqElim** can have only one possible result: the given value, considered as having the output type. Then, we abstract a singleton set. This means we can treat each application of **eqElim** to $?$ as an application to **Refl** $?$ $?$. This principle holds for any single-constructor inductive type.

With only functions, we needed to return $?$ any time we applied a dynamically-typed function. However, with eliminators, we are always structurally decreasing on the value being eliminated. For **natElim**, we can eliminate a $?$ -typed value provided it is 0 or **Succ** u , but otherwise we must produce $?$ for substitution to be total and type-preserving. Other types are handled similarly.

Runtime Semantics. The semantics are straightforward. Eliminations are handled as with hereditary substitution: eliminating $?$ produces $?$, except with **Eq**, where $?$ behaves like **Refl** $?$ $?$ when eliminating. When applying eliminators or constructors, evidence is composed as with functions.

One advantage of GDTL is that the meet operator on evidence allows definitional equality checks to be moved to runtime. Thus, if we write `Refl ? ?`, but use it at type `Eq A x y`, then it is ascribed with evidence $\langle \text{Eq } A \ x \ y \rangle$. If we ever use this proof to transform a value of type `P x` into one of type `P y`, the meet operation on the evidence ensures that the result actually has type `P y`.

Returning to the head' function from Section 2.3, in `head' Nat 0 ? ? staticNil`, the second `?` is ascribed with the evidence $\langle \text{Eq } \text{Nat } 0 \ (\text{Succ } ?) \rangle$. The call to `rewrite` using this proof tries to convert a `Vec` of length 0 into one of length 1, which adds the evidence $\langle \text{Eq } \text{Nat } 0 \ 1 \rangle$ to our proof term. Evaluation tries to compose the layered evidence, but fails with the rule `STEPASCRFAIL`, since they cannot be composed.

9 RELATED WORK

SDTL. The static dependently-typed language SDTL, from which GDTL is derived, incorporates many features and techniques from the literature. The core of the language is very similar to that of CC_ω [Coquand and Huet 1988], albeit without an impredicative `Prop` sort. The core language of Idris [Brady 2013], TT, also features cumulative universes with a single syntactic category for terms and types. Our use of canonical forms draws heavily from work on the Logical Framework (LF) [Harper et al. 1993; Harper and Licata 2007]. The bidirectional type system we adopt is inspired by the tutorial of Löh et al. [2010]. Our formulation of hereditary substitution [Pfenning 2008; Watkins et al. 2003] in SDTL is largely drawn from that of Harper and Licata [2007], particularly the type-outputting judgment for substitution on atomic forms, and the treatment of the variable type as an extrinsic argument.

Mixing Dependent Types and Non-termination. Dependently-typed languages that admit non-termination either give up on logical consistency altogether (Ω mega [Sheard and Linger 2007], Haskell), or isolate a sublanguage of pure terminating expressions. This separation can be either enforced through the type system and/or a termination checker (Aura [Jia et al. 2008], F★ [Swamy et al. 2016], Idris [Brady 2013]), or through a strict syntactic separation (Dependent ML [Xi and Pfenning 1999], ATS [Chen and Xi 2005]). The design space is very wide, reflecting a variety of sensible tradeoffs between expressiveness, guarantees, and flexibility.

The Zombie language [Casinghino et al. 2014; Sjöberg et al. 2012] implements a flexible combination of programming and proving. The language is defined as a the programmatic fragment that ensures type safety but not termination, and a logical fragment (a syntactic subset of the programmatic one) that guarantees logical consistency. Programmers must declare in which fragment a given definition lives, but mobile types and cross-fragment case expressions allow interactions between the fragments. Zombie embodies a different tradeoff from GDTL: while the logical fragment is consistent as a logic, typechecking may diverge due to normalization of terms from the programmatic fragment. In contrast, GDTL eschews logical consistency as soon as imprecision is introduced (with `?`), but approximate normalization ensures that typechecking terminates.

In general, gradual dependent types as provided in GDTL can be interpreted as a smooth, tight integration of such a two-fragment approach. Indeed, the subset of GDTL that is fully precise corresponds to SDTL, which is consistent as a logic. However, in the gradual setting, the fragment separation is fluid: it is driven by the precision of types and terms, which is free to evolve at an arbitrarily fine level of granularity. Also, the mentioned approaches are typically not concerned with accommodating the flexibility of a fully dynamically-typed language.

Mixing Dependent Types and Simple Types. Several approaches have explored how soundly combine dependent types with non-dependently typed components. Ou et al. [2004] support a two-fragment language, where runtime checks at the boundary ensure that dependently-typed properties hold. The approach is limited to properties that are expressible as boolean-valued

functions. [Tanter and Tabareau \[2015\]](#) develop a cast framework for subset types in Coq, allowing one to assert a term of type A to have the subset type $\{a : A \mid P\ a\}$ for any decidable (or checkable) property P . They use an axiom to represent cast errors. [Osera et al. \[2012\]](#) present *dependent interoperability* as a multi-language approach to combine both typing disciplines, mediated by runtime checks. Building on the subset cast framework (reformulated in a monadic setting instead of using axiomatic errors), [Dagand et al. \[2016, 2018\]](#) revisit dependent interoperability in Coq via type-theoretic Galois connections that allow automatic lifting of higher-order programs. Dependent interoperability allows exploiting the connection between pre-existing types, such as **Vec** and **List**, imposing the overhead of data structure conversions. The fact that **List** is a less precise structure than **Vec** is therefore defined *a posteriori*. In contrast, in GDTL, one can simply define **List** A as an alias for **Vec** A ?, thereby avoiding the need for deep structural conversions.

The work of [Lehmann and Tanter \[2017\]](#) on gradual refinement types includes some form of dependency in types. Gradual refinement types range from simple types to logically-refined types, i.e. subset types where the refinement is drawn from an SMT-decidable logic. Imprecise logical formulae in a function type can refer to arguments and variables in context. This kind of value dependency is less expressive than the dependent type system considered here. Furthermore, GDTL is the first gradual language to allow ? to be used in both term and type position, and to fully embed the untyped lambda calculus.

Programming with Holes. Finally, we observe that using ? in place of proof terms in GDTL is related to the concept of *holes* in dependently-typed languages. Idris [[Brady 2013](#)] and Agda [[Norell 2009](#)] both allow typechecking of programs with typed holes. The main difference between ? and holes in these languages is that applying a hole to a value results in a stuck term, while in GDTL, applying ? to a value produces another ?.

Recently, [Omar et al. \[2019\]](#) describe Hazelnut, a language and programming system with *typed holes* that fully supports evaluation in presence of holes, including reduction *around* holes. The approach is based on Contextual Modal Type Theory [[Nanevski et al. 2008](#)]. It would be interesting to study whether the dependently-typed version of CMTT [[Pientka and Dunfield 2008](#)] could be combined with the evaluation approach of Hazelnut, and the IDE support, in order to provide a rich programming experience with gradual dependent types.

10 CONCLUSION

GDTL represents a glimpse of the challenging and potentially large design space induced by combining dependent types and gradual typing. Specifically, this work proposes approximate normalization as a novel technique for designing gradual dependently-typed languages, in a way that ensures decidable typechecking and naturally satisfies the gradual guarantees.

Currently, GDTL lacks a number of features required of a practical dependently-typed programming language. While we have addressed the most pressing issue of supporting inductive types in Section 8, the metatheory of this extension, in particular the proof of strong normalization, is future work. It might also be interesting to consider pattern matching as the primitive notion for eliminating inductives, as in Agda, instead of elimination principles as in Coq; the equalities implied by dependent matches could be turned into runtime checks for gradually-typed values.

Future work includes supporting implicit arguments and higher-order unification, blame tracking [[Wadler and Findler 2009](#)], and efficient runtime semantics with erasure of computationally-irrelevant arguments [[Brady et al. 2003](#)]. Approximate normalization might be made more precise by exploiting *termination contracts* [[Nguyễn et al. 2019](#)].

Acknowledgments. We thank the anonymous reviewers for their constructive feedback.

REFERENCES

- Felipe Bañados Schwerter, Ronald Garcia, and Éric Tanter. 2016. Gradual Type-and-Effect Systems. *Journal of Functional Programming* 26 (Sept. 2016), 19:1–19:69.
- Yves Bertot and Pierre Castéran. 2004. *Interactive Theorem Proving and Program Development*. Springer-Verlag.
- Edwin Brady. 2013. Idris, a general-purpose dependently typed programming language: Design and implementation. *Journal of Functional Programming* 23, 5 (2013), 552–593. <https://doi.org/10.1017/S095679681300018X>
- Edwin Brady, Conor McBride, and James McKinna. 2003. Inductive Families Need Not Store Their Indices. In *Types for Proofs and Programs, International Workshop, TYPES 2003, Torino, Italy, April 30 - May 4, 2003, Revised Selected Papers (Lecture Notes in Computer Science)*, Stefano Berardi, Mario Coppo, and Ferruccio Damiani (Eds.), Vol. 3085. Springer, 115–129. https://doi.org/10.1007/978-3-540-24849-1_8
- Chris Casinghino, Vilhelm Sjöberg, and Stephanie Weirich. 2014. Combining Proofs and Programs in a Dependently Typed Language. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '14)*. ACM, New York, NY, USA, 33–45. <https://doi.org/10.1145/2535838.2535883>
- Ilano Cervesato and Frank Pfenning. 2003. A Linear Spine Calculus. *Journal of Logic and Computation* 13, 5 (2003), 639–688. <https://doi.org/10.1093/logcom/13.5.639> arXiv:<http://logcom.oxfordjournals.org/content/13/5/639.full.pdf+html>
- Chiyan Chen and Hongwei Xi. 2005. Combining Programming with Theorem Proving. In *Proceedings of the Tenth ACM SIGPLAN International Conference on Functional Programming (ICFP '05)*. ACM, New York, NY, USA, 66–77. <https://doi.org/10.1145/1086365.1086375>
- Adam Chlipala. 2013. *Certified Programming with Dependent Types: A Pragmatic Introduction to the Coq Proof Assistant*. The MIT Press.
- T. Coquand. 1986. *An analysis of Girard's paradox*. Technical Report RR-0531. INRIA. <https://hal.inria.fr/inria-00076023>
- Thierry Coquand and Gérard Huet. 1988. The calculus of constructions. *Information and Computation* 76, 2 (1988), 95 – 120. [https://doi.org/10.1016/0890-5401\(88\)90005-3](https://doi.org/10.1016/0890-5401(88)90005-3)
- Pierre-Evariste Dagand, Nicolas Tabareau, and Éric Tanter. 2016. Partial Type Equivalences for Verified Dependent Interoperability. In *Proceedings of the 21st ACM SIGPLAN Conference on Functional Programming (ICFP 2016)*. ACM Press, Nara, Japan, 298–310.
- Pierre-Évariste Dagand, Nicolas Tabareau, and Éric Tanter. 2018. Foundations of dependent interoperability. *Journal of Functional Programming* 28 (2018), e9. <https://doi.org/10.1017/S0956796818000011>
- Nachum Dershowitz and Zohar Manna. 1979. Proving termination with multiset orderings. In *Automata, Languages and Programming*, Hermann A. Maurer (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 188–202.
- Harley Eades and Aaron Stump. 2010. Hereditary substitution for stratified system F. In *International Workshop on Proof-Search in Type Theories, PSTT*, Vol. 10.
- Richard A. Eisenberg. 2016. Dependent Types in Haskell: Theory and Practice. arXiv:[cs.PL/1610.07978](https://arxiv.org/abs/1610.07978)
- Joseph Eremondi. 2019. Github Repository: GDTL-artifact. <https://github.com/JoeyEremondi/GDTL-artifact/>.
- Joseph Eremondi, Éric Tanter, and Ronald Garcia. 2019. Approximate Normalization for Gradual Dependent Types. arXiv:[1906.06469](https://arxiv.org/abs/1906.06469)
- Luminous Fennell and Peter Thiemann. 2013. Gradual Security Typing with References. In *Proceedings of the 26th Computer Security Foundations Symposium (CSF)*. 224–239.
- Ronald Garcia, Alison M. Clark, and Éric Tanter. 2016. Abstracting Gradual Typing. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '16)*. ACM, New York, NY, USA, 429–442. <https://doi.org/10.1145/2837614.2837670>
- Jean-Yves Girard. 1972. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. Ph.D. Dissertation. PhD thesis, Université Paris VII.
- Ben Greenman and Matthias Felleisen. 2018. A Spectrum of Type Soundness and Performance. *Proc. ACM Program. Lang.* 2, ICFP, Article 71 (July 2018), 32 pages. <https://doi.org/10.1145/3236766>
- Robert Harper, Furio Honsell, and Gordon Plotkin. 1993. A Framework for Defining Logics. *J. ACM* 40, 1 (Jan. 1993), 143–184. <https://doi.org/10.1145/138027.138060>
- Robert Harper and Daniel R. Licata. 2007. Mechanizing metatheory in a logical framework. *Journal of Functional Programming* 17, 4-5 (2007), 613–673. <https://doi.org/10.1017/S0956796807006430>
- Atsushi Igarashi, Peter Thiemann, Vasco T. Vasconcelos, and Philip Wadler. 2017. Gradual Session Types. *Proceedings of the ACM on Programming Languages* 1, ICFP (Sept. 2017), 38:1–38:28.
- Limin Jia, Jeffrey A. Vaughan, Karl Mazurak, Jianzhou Zhao, Luke Zarko, Joseph Schorr, and Steve Zdancewic. 2008. AURA: a programming language for authorization and audit. In *Proceeding of the 13th ACM SIGPLAN international conference on Functional programming, ICFP 2008, Victoria, BC, Canada, September 20-28, 2008*, James Hook and Peter Thiemann (Eds.). ACM, 27–38. <https://doi.org/10.1145/1411204.1411212>
- Nico Lehmann and Éric Tanter. 2017. Gradual Refinement Types. In *Proceedings of the 44th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2017)*. ACM Press, Paris, France, 775–788.

- Andres Löb, Conor McBride, and Wouter Swierstra. 2010. A Tutorial Implementation of a Dependently Typed Lambda Calculus. *Fundam. Inf.* 102, 2 (April 2010), 177–207. <https://doi.org/10.3233/FI-2010-304>
- Cyprien Mangin and Matthieu Sozeau. 2015. Equations for Hereditary Substitution in Leivant’s Predicative System F: A Case Study. *CoRR* abs/1508.00455 (2015). arXiv:1508.00455 <http://arxiv.org/abs/1508.00455>
- Aleksandar Nanevski, Frank Pfenning, and Brigitte Pientka. 2008. Contextual Modal Type Theory. *ACM Trans. Comput. Logic* 9, 3, Article 23 (June 2008), 49 pages. <https://doi.org/10.1145/1352582.1352591>
- Phúc C. Nguyễn, Thomas Gilray, Sam Tobin-Hochstadt, and David Van Horn. 2019. Size-change Termination As a Contract: Dynamically and Statically Enforcing Termination for Higher-order Programs. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2019)*. ACM, New York, NY, USA, 845–859. <https://doi.org/10.1145/3314221.3314643>
- Ulf Norell. 2009. Dependently Typed Programming in Agda. In *Proceedings of the 4th International Workshop on Types in Language Design and Implementation (TLDI '09)*. ACM, New York, NY, USA, 1–2. <https://doi.org/10.1145/1481861.1481862>
- Cyrus Omar, Ian Voysey, Ravi Chugh, and Matthew A. Hammer. 2019. Live functional programming with typed holes. *PACMPL* 3, POPL (2019), 14:1–14:32. <https://dl.acm.org/citation.cfm?id=3290327>
- Peter-Michael Osera, Vilhelm Sjöberg, and Steve Zdancewic. 2012. Dependent Interoperability. In *Proceedings of the Sixth Workshop on Programming Languages Meets Program Verification (PLPV '12)*. ACM, New York, NY, USA, 3–14. <https://doi.org/10.1145/2103776.2103779>
- Xinming Ou, Gang Tan, Yitzhak Mandelbaum, and David Walker. 2004. Dynamic Typing with Dependent Types. In *Exploring New Frontiers of Theoretical Informatics*, Jean-Jacques Levy, Ernst W. Mayr, and John C. Mitchell (Eds.). Springer US, Boston, MA, 437–450.
- Frank Pfenning. 2008. Church and Curry: Combining intrinsic and extrinsic typing. In *Reasoning in Simple Type Theory – Festschrift in Honor of Peter B. Andrews on His 70th Birthday (Studies in Logic, Mathematical Logic and Foundations)*, Christoph Benzmüller, Chad Brown, Jörg Siekmann, and Richard Statman (Eds.). College Publications. <http://www.collegepublications.co.uk/logic/mlf/?00010>
- Brigitte Pientka and Joshua Dunfield. 2008. Programming with Proofs and Explicit Contexts. In *Proceedings of the 10th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP '08)*. ACM, New York, NY, USA, 163–173. <https://doi.org/10.1145/1389449.1389469>
- Benjamin C. Pierce and David N. Turner. 2000. Local Type Inference. *ACM Trans. Program. Lang. Syst.* 22, 1 (Jan. 2000), 1–44. <https://doi.org/10.1145/345099.345100>
- Ilya Sergey and Dave Clarke. 2012. Gradual Ownership Types. In *Proceedings of the 21st European Symposium on Programming Languages and Systems (ESOP 2012) (Lecture Notes in Computer Science)*, Helmut Seidl (Ed.), Vol. 7211. Springer-Verlag, Tallinn, Estonia, 579–599.
- Tim Sheard and Nathan Linger. 2007. Programming in Omega. In *Central European Functional Programming School, Second Summer School, CEFP 2007, Cluj-Napoca, Romania, June 23-30, 2007, Revised Selected Lectures (Lecture Notes in Computer Science)*, Zoltán Horváth, Rinus Plasmeijer, Anna Soós, and Viktória Zsóka (Eds.), Vol. 5161. Springer, 158–227. https://doi.org/10.1007/978-3-540-88059-2_5
- Jeremy G. Siek and Walid Taha. 2006. Gradual Typing for Functional Languages. In *Scheme and Functional Programming Workshop*. 81–92.
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. 2015. Refined Criteria for Gradual Typing. In *1st Summit on Advances in Programming Languages (SNAPL 2015) (Leibniz International Proceedings in Informatics (LIPIcs))*, Thomas Ball, Rastislav Bodik, Shriram Krishnamurthi, Benjamin S. Lerner, and Greg Morrisett (Eds.), Vol. 32. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 274–293. <https://doi.org/10.4230/LIPIcs.SNAPL.2015.274>
- Jeremy G. Siek and Philip Wadler. 2010. Threesomes, with and Without Blame. In *Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '10)*. ACM, New York, NY, USA, 365–376. <https://doi.org/10.1145/1706299.1706342>
- Vilhelm Sjöberg, Chris Casinghino, Ki Yung Ahn, Nathan Collins, Harley D. Eades III, Peng Fu, Garrin Kimmell, Tim Sheard, Aaron Stump, and Stephanie Weirich. 2012. Irrelevance, Heterogeneous Equality, and Call-by-value Dependent Type Systems. In *Proceedings Fourth Workshop on Mathematically Structured Functional Programming*, Tallinn, Estonia, 25 March 2012 (*Electronic Proceedings in Theoretical Computer Science*), James Chapman and Paul Blain Levy (Eds.), Vol. 76. Open Publishing Association, 112–162. <https://doi.org/10.4204/EPTCS.76.9>
- Aaron Stump. 2017. The calculus of dependent lambda eliminations. *Journal of Functional Programming* 27 (2017), e14. <https://doi.org/10.1017/S0956796817000053>
- Nikhil Swamy, Cătălin Hrițcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean-Karim Zinzindohoue, and Santiago Zanella-Béguelin. 2016. Dependent Types and Multi-Monadic Effects in F*. In *43rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. ACM, 256–270. <https://www.fstar-lang.org/papers/mumon/>

- Éric Tanter and Nicolas Tabareau. 2015. Gradual Certified Programming in Coq. In *Proceedings of the 11th Symposium on Dynamic Languages (DLS 2015)*. ACM, New York, NY, USA, 26–40. <https://doi.org/10.1145/2816707.2816710>
- Matias Toro, Ronald Garcia, and Éric Tanter. 2018a. Type-Driven Gradual Security with References. *ACM Transactions on Programming Languages and Systems* 40, 4 (Nov. 2018), 16:1–16:55.
- Matias Toro, Elizabeth Labrada, and Éric Tanter. 2018b. Gradual Parametricity, Revisited. arXiv:[cs.PL/1807.04596](https://arxiv.org/abs/cs.PL/1807.04596)
- Michael M. Vitousek, Cameron Swords, and Jeremy G. Siek. 2017. Big Types in Little Runtime: Open-world Soundness and Collaborative Blame for Gradual Type Systems. *SIGPLAN Not.* 52, 1 (Jan. 2017), 762–774. <https://doi.org/10.1145/3093333.3009849>
- Philip Wadler and Robert Bruce Findler. 2009. Well-Typed Programs Can't Be Blamed. In *Programming Languages and Systems*, Giuseppe Castagna (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1–16.
- Kevin Watkins, Iliano Cervesato, Frank Pfenning, and David Walker. 2003. *A Concurrent Logical Framework I: Judgments and Properties*. Technical Report CMU-CS-02-101. <https://www.cs.cmu.edu/~fp/papers/CMU-CS-02-101.pdf>.
- Hongwei Xi and Frank Pfenning. 1999. Dependent Types in Practical Programming. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '99)*. ACM, New York, NY, USA, 214–227. <https://doi.org/10.1145/292540.292560>